AD-A118 828 H H AEROSPACE DESIGN CO INC BEDFORD MA F/G 9/2
A TECHNICAL ANALYSIS OF REHOSTING THE NATIONAL AIRSPACE SYSTEM --ETC(U)
OCT 81 K BRITTING, A COCANOWER, A GUPTA DT-FA01-80-Y-10595
UNCLASSIFIED HHA-81-4 FAA-AP-82-2 NL

1 of 2
AD
A118828

October 1981

Technical Report Documentation Page

1. Report No. FAA-AP-82-2	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle: A Technical Analysis of Rehosting the National Airspace System Software		5. Report Date: October 1981
		6. Performing Organization Code
7. Author(s): K. Britting, A. Cocanower, A. Gupta, M. Hsu, F. Marcus, J. Oiesen, H. Toong		8. Performing Organization Report No. HHA-81-4
9. Performing Organization Name and Address: HH Aerospace Design Company, Inc. Civil Air Terminal Bedford, MA 01730		10. Work Unit No. (TRAIS)
		11. Contract or Grant No. DT-FA01-80-Y-10595
12. Sponsoring Agency Name and Address: Federal Aviation Administration Advanced Automation Program Office 800 Independence Avenue, S.W. Washington, D.C. 20591		13. Type of Report and Period Covered: Final Report November 1980-August 1981
		14. Sponsoring Agency Code: AAP
15. Supplementary Notes		

16. Abstract

The Federal Aviation Administration (FAA) is planning to replace within the next ten years the computers used to provide en route air traffic control services; in carrying out this replacement there are many different strategies the FAA could follow. The purpose of this report is to study the strategy known as rehosting the National Airspace System (NAS) software on instruction-compatible machines. The idea is that the current computers (and associated peripherals) would be replaced by modern hardware that executes the same machine-language instructions. The current NAS software would be changed only insofar as proves necessary for the software to run on the new machines; these changes to the software are expected to be minor.

The rehosting strategy is evaluated in seven areas. First, how reliable is the system? Second, how well will the system perform under expected workloads? Third, how serious are the technical obstacles to adapting the software to run on the new machines? Fourth, what would the new system cost? Fifth, what problems would be encountered during the transition to the new system. Sixth, how quickly could the system be procured? Seventh, how well adapted is the system to future growth.

The conclusion is that the rehost strategy is technically feasible, but there is some uncertainty about what this strategy would cost and how long the procurement process would take.

17. Key Words: Air Traffic Control, Reliability, Computer Performance, Transition, IBM 9020 Computers	18. Distribution Statement

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified		

Form DOT F 1700.7 (8-72)   Reproduction of completed page authorized

NOTICE

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assures no liability for its contents or use thereof.

NOTICE

The United States Government does not endorse products or manufacturers. Trade or manufacturer's names appear herein solely because they are considered essential to the object of this report.

METRIC CONVERSION FACTORS

Approximate Conversions to Metric Measures

Symbol	When You Know	Multiply by	To Find	Symbol
		LENGTH		
in	inches	*2.5	centimeters	cm
ft	feet	30	centimeters	cm
yd	yards	0.9	meters	m
mi	miles	1.6	kilometers	km
		AREA		
in^2	square inches	6.5	square centimeters	cm^2
ft^2	square feet	0.09	square meters	m^2
yd^2	square yards	0.8	square meters	m^2
mi^2	square miles	2.6	square kilometers	km^2
	acres	0.4	hectares	ha
		MASS (weight)		
oz	ounces	28	grams	g
lb	pounds	0.45	kilograms	kg
	short tons (2000 lb)	0.9	tonnes	t
		VOLUME		
tsp	teaspoons	5	milliliters	ml
Tbsp	tablespoons	15	milliliters	ml
fl oz	fluid ounces	30	milliliters	ml
c	cups	0.24	liters	l
pt	pints	0.47	liters	l
qt	quarts	0.95	liters	l
gal	gallons	3.8	liters	l
ft^3	cubic feet	0.03	cubic meters	m^3
yd^3	cubic yards	0.76	cubic meters	m^3
		TEMPERATURE (exact)		
°F	Fahrenheit temperature	5/9 (after subtracting 32)	Celsius temperature	°C

*1 in ≡ 2.54 (exactly). For other exact conversions and more detailed tables, see NBS Misc. Publ. 286, Units of Weights and Measures, Price $2.25, SD Catalog No. C13.10:286.

Approximate Conversions from Metric Measures

Symbol	When You Know	Multiply by	To Find	Symbol
		LENGTH		
mm	millimeters	0.04	inches	in
cm	centimeters	0.4	inches	in
m	meters	3.3	feet	ft
m	meters	1.1	yards	yd
km	kilometers	0.6	miles	mi
		AREA		
cm^2	square centimeters	0.16	square inches	in^2
m^2	square meters	1.2	square yards	yd^2
km^2	square kilometers	0.4	square miles	mi^2
ha	hectares (10,000 m^2)	2.5	acres	
		MASS (weight)		
g	grams	0.035	ounces	oz
kg	kilograms	2.2	pounds	lb
t	tonnes (1000 kg)	1.1	short tons	
		VOLUME		
ml	milliliters	0.03	fluid ounces	fl oz
l	liters	2.1	pints	pt
l	liters	1.06	quarts	qt
l	liters	0.26	gallons	gal
m^3	cubic meters	35	cubic feet	ft^3
m^3	cubic meters	1.3	cubic yards	yd^3
		TEMPERATURE (exact)		
°C	Celsius temperature	9/5 (then add 32)	Fahrenheit temperature	°F

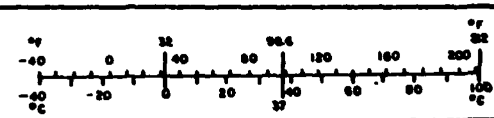

PREFACE

We would like to thank the people and organizations that have given us information and assistance during this study. The people who have taken the time to talk to us are unfortunately too numerous to list, so we will instead mention their organizations. At FAA headquarters in Washington we talked to personnel in the Office of Systems Engineering Management, the Airway Facilities (AF) Service, the Air Traffic (AT) Service, the Office of Personnel and Training, the Office of the Associate Administrator for Engineering and Development, and the Office of Aviation Systems Plans. At the FAA Technical Center in Atlantic City we talked to people in R&D and ACT-700 as well as to additional people in AF and AT. At the Mike Monroney Aeronautical Center in Oklahoma City we talked to personnel in the Depot, the Academy, and the Budget Division. Amdahl and IBM kindly hosted us and talked at length. The Analytic Sciences Corporation, which is conducting a related study, assisted us at many points. Finally, we would like to single out for thanks Dr. Andres Zellweger, the technical monitor, who provided guidance, advice, and encouragement throughout this study.

The contributors to this report are Kenneth R. Britting, Alfred B. Cocanower, Amar Gupta, Mei Hsu, Fred J. Marcus, James F. Oiesen, and Hoo-min D. Toong.

CONTENTS

Appendix

EXECUTIVE SUMMARY

The Federal Aviation Administration (FAA) is planning to replace by 1990 the computer systems used to provide en route air traffic control services. If, however, air traffic grows so that the demand on these computers exceeds their capacity before they are replaced, then it will be necessary to either restrict air traffic or to adopt some interim system designed to stretch the life of the current system by a few years. This report is one in a series that examines the advantages and disadvantages of the potential interim systems in order to provide the information needed by FAA decision-makers.

The purpose of this report is to discuss the interim system achieved by rehosting the National Airspace System (NAS) software on instruction-compatible machines. That is, the current computers, collectively referred to as IBM 9020's, at each air route traffic control center (ARTCC) would be replaced by modern machines that execute nearly the same instruction set. This rehost system would use the NAS software currently used, with this software only changed insofar as necessary for it to run on the new machines. The advantage sometimes claimed for this system is that it would allow the FAA to increase the capacity and reliability of the en route computer systems while avoiding the expense and risk of completely new software.

The rehost system would replace the current computers (including the display channel computers), tapes, and disks. The peripheral adapter modules and the controller suites up to and including the display generators would not be replaced. The heart of the rehost system would be two mainframes. One mainframe would handle the processing now done by the central computer complex (CCC) and by the display channel; the other mainframe would be standing by ready to take over the processing if the first mainframe fails. The back-up mainframe would be able to carry out ancillary processing tasks (e.g., analyzing performance data, providing training simulations) while standing by.

For concreteness this report analyzes one specific rehost system. Several variants have been sugested, but these are not considered in detail since it is expected that they would not significantly alter the analysis.

The information about the rehost system that is relevant to the FAA's decision on whether this system should be procured is presented under the headings of cost, schedule, reliability, performance, technical issues, transition, and growth potential.

Cost. Using the current cost of providing en route air traffic control services as a baseline, the change in this cost that would result from rehosting is estimated. Seven categories of cost are considered. First, the cost of developing and initially testing the software is estimated to be $5.8 million. This cost covers the needed modifications to the on-line software, the support software, and the virtual machine monitor. Second, the cost of acquiring the hardware for 23 sites is estimated to be either $123.8 million if Amdahl 470/V7's are purchased or $175.1 million if IBM 3033U's are purchased. Since the V7 and the 3033U are judged to be the two mainframes that are best suited to rehosting, the cost calculation is carried out for both. These figures include the cost of mainframes, tape units, disk units, other peripherals, and the special hardware that is needed. Third, the cost of testing the complete system at the FAA Technical Center and at the first en route center is estimated to be $6.3 million. This cost covers the testing that is necessary to bring the system to the point where it is ready to be routinely deployed. Fourth, the cost incurred during the transition period is estimated to be $36.9 million. This figure covers the cost of remodeling the centers, the cost of the extra personnel needed during transition, and the cost of developing courses on the new system and teaching them to FAA personnel. Fifth, the initial cost of spare parts and documentation is expected to be either $26.7 million if V7's are purchased or $37.8 million if 3033U's are purchased. Sixth, because of the greater reliability of the rehost system, there will be a saving in the cost of spare parts and maintenance personnel. After an initial shakedown period, the rehost system would save an estimated $9.3 million per year. Seventh, the FAA administrative cost over the six year program is estimated to be $41.2 million. This covers program planning, management, and review.

Table ES-1 summarizes all of the front-end costs that are incurred to get the rehost system operational at all sites. This table shows each category of cost and indicates where it would be incurred; HQ, TC, and AC stand for FAA headquarters, the FAA Technical Center, and the FAA Aeronautical Center, respectively. The initial cost that would be incurred if rehosting were adopted is estimated to be either $241.0 million if V7's are procured or $303.4 million if 3033U's are procured. All costs are in 1981 dollars. Table ES-2 shows how the annual saving in the maintenance cost would begin at about a half million dollars in the first year [illegible]hich a system is installed and would gradually rise to $9.3 million per [illegible]r.

The main cost of rehosting is seen to be the hardware acquisi[illegible]st, which is more than half the initial cost. Since most of the hardw[illegible] acquired is off-the-shelf equipment, there is relatively little uncertainty about this cost. Because of the uniqueness of the rehost problem, there is considerable uncertainty about the administrative costs, the software cost, the spare parts cost (which might well be overestimated), and the transition cost (which might well be underestimated).

These cost estimates assume that there is replacement at all 20 ARTCC's. It is possible, however, that the rehost system would be installed only at those centers that faced an imminent capacity problem. This partial replacement would have the advantage of cutting down the cost considerably; this saving largely results from avoiding the hardware acquisition cost, which is the main cost, at those centers at which there is no capacity problem. It is estimated that if V7's are procured, then the initial cost of rehosting would be $107.0 million if there were replacement at five centers and $151.7 million if there were replacement at ten centers, compared to the cost of $241.0 million if there were replacement at all twenty centers. A disadvantage of partial replacement is that support would be complicated since two entirely different systems would be in the field.

Schedule. Once the FAA decided to rehost and issued a request for proposals (RFP), the steps in the procurement and the estimated length of each step would be:

TABLE ES-1: INITIAL COSTS INCURRED BY REHOSTING (millions of dollars)

	Site				
	HQ	TC	AC	ARTCC's	Total
Software	5.8				5.8
Hardware					
Engineering	0.3				0.3
Acquisition					
V7		10.8	5.4	107.6	123.8
3033U		15.2	7.6	152.3	175.1
Testing		3.8		2.5	6.3
Maintenance					
Initial cost					
V7			0.9	25.8	26.7
3033U			1.2	36.6	37.8
Transition Cost					
Remodeling		2.0	1.0	20.0	23.0
Extra personnel				4.0	4.0
Developing courses			1.8		1.8
Teaching courses		0.6	0.1	7.4	8.1
Program management and Support	41.2				41.2
Total					
V7					241.0
3033U					303.4

TABLE ES-2: ANNUAL MAINTENANCE COST SAVING PROVIDED BY REHOSTING

Year	Saving (millions)
1	$0.506
2	2.387
3	5.135
4	7.883
5 and after	9.257

- industry prepares proposals (3 months);
- FAA evaluates the proposals and awards a contract (6 months);
- contractor develops hardware and software (21 months);
- FAA and contractor test a system at the FAA Technical Center (9 months);
- FAA and contractor test a system at the first field site (6 months);
- contractor installs systems at the remaining sites (24 months).

Therefore, from the time that an RFP is issued to the time that the system is operational at the first field site, there is an elapsed time of 45 months (3 years, 9 months). From the time an RFP is issued until the system is operational at all sites, 69 months (5 years, 9 months) elapses. This means that if an RFP is issued on 1 July 1982, the first system will be operational at an en route center on 1 April 1986, and the system will be operational at all centers on 1 April 1988.

One suggested rehosting approach differs from the approach considered here by retaining the disk and tape drives, by making fewer software changes, and by incurring a greater processing overhead. This approach would reduce the development time by an estimated 12 to 15 months and would reduce the cost by an estimated $25 million; there would, however, be greater uncertainty over the schedule and cost.

It has also been suggested that the mainframes should be leased rather than purchased to shorten the procurement cycle and to save money; this approach has three problems. First, since developing the system rather than acquiring the mainframes is the bottleneck in the procurement, leasing would not speed the procurement. Second, since three years is typically the break-even point for a lease and since these computers would probably be in place for more than three years, leasing would probably end up costing more rather than less. Third, the user typically is not allowed to maintain leased computers, and this would interfere with the FAA's providing the type of maintenance required by air traffic control. (It should be remarked, however, that because of changes in technology the FAA would probably do less of the maintenance for the rehost system; for example, the FAA might use the manufacturer's remote diagnostic services.)

Reliability. The FAA's goal is to have a system with extremely high availability, i.e., a system that supports air traffic control with minimal interruptions in service. The types of failures that can beset the system are hardware, software, and personnel failures.

In discussing hardware failures it is essential to distinguish between a component failure and a system failure. For example, a 9020D has three compute elements (CE's), two of which are active and one of which is redundant. If one of the active CE's fails, the system is automatically reconfigured so that the redundant CE is made active; with two active CE's, there is no system failure even though a component has failed. A system failure occurs only if, before the CE that failed is repaired or replaced, one of the other CE's fails. Therefore, because of the redundancy built into the system, a single component failure does not cause a system failure; a system failure only results when the number of failed components exceeds the number of redundant components. Redundancy, then, can lessen but not completely eliminate system failures. The rehost system would decrease the time it takes to reconfigure the system when a component fails; the rehost system would also decrease the frequency of system failures. These improvements would reduce the uncertainty that controllers now have about the ability of the system to quickly and completely recover from a failure.

The relative availability of the 9020 and rehost systems is studied by combining information about the redundancy built into each system with information about the mean time between failure (MTBF) and the mean repair time of each component. A system failure occurs in the 9020D if at least one of the following conditions is violated:

- at least 2 of the 3 compute elements are working;
- at least 5 of the 6 storage elements are working;
- at least 2 of the 3 input/output control elements are working;
- at least 2 of the 3 tape control units are working;
- at least 2 of the 3 disk control units are working.

In the rehost system a mainframe is said to contain a CPU, a memory, and 12 channels divided into 6 pairs. A mainframe is working if the CPU is working, if the memory is working, and if at least one channel in each pair is working. A system failure occurs in the rehost system if at least one of the following conditions is violated:

- at least 1 of the 2 mainframes is working;
- at least 1 of the 2 tape control units are working;
- at least 1 of the 2 disk control units are working.

Once assumptions are made about the MTBF's of each component and the mean time to repair, the system availability and MTBF are calculated; these results are shown in Table ES-3 for the rehost system and for a system with a 9020D in the CCC and a 9020E in the display channel. The mean time between system failures is estimated to be 2905 days for the rehost hardware and 1226 days for the 9020D/9020E hardware. This greater reliability of the rehost hardware stems largely from its being configured in parallel; that is, if either mainframe fails, the system does not fail since the remaining mainframe can carry the entire load. The 9020D/9020E system, in contrast, is configured in series; that is, the system operates only if both the 9020D and the 9020E operate. There is a good deal of uncertainty about the accuracy of these estimates because the data available for determining the component MTBF's and repair times were sketchy. Therefore, a sensitivity analysis was carried out, and it was found that the rehost hardware retained

TABLE ES-3: ESTIMATED AVAILABILITY AND MTBF OF SYSTEM HARDWARE

System	Availability	MTBF (days)
9020D/9020E	0.99998301	1226
Rehost	0.99999283	2905

its reliability advantage for alternate values of the component MTBF's and repair times. In sum, while the absolute numbers might be questioned, the conclusion is that the rehost system does exhibit greater hardware reliability because the results are so lopsided in favor of the rehost system and because of the persistence of this finding throughout the sensitivity analysis.

Now turn to the topic of software reliability. The rehost system has three major software components: the NAS application software, the NAS monitor, and the virtual machine monitor. During the testing phase it is expected that new problems would arise with the NAS monitor and application code, but by the time the rehost system is operational it is expected that these two components will return to their present level of reliability. In fact, because the NAS software will all be memory-resident, the swapping of code in and out of main memory will be eliminated. Because swapping and table size limitations are a significant source of software failures, the NAS software can be expected to be more reliable under rehosting. The virtual machine monitor would be a source of new software failures. Therefore, under rehosting there will be a decrease in failures because swapping is eliminated and an increase in failures because of the virtual machine monitor. It is impossible to quantify the net effect on software reliability, but it can be concluded that, at worst, the rehost system will have only slightly lower software reliability.

In order to estimate the overall availability and MTBF of the system where both hardware and software failures are considered, the assumptions made about the hardware are supplemented with tentative assumptions about the frequency and duration of software failures. Table ES-4 shows the

TABLE ES-4: ESTIMATED AVAILABILITY AND MTBF OF SYSTEM HARDWARE AND SOFTWARE

System	Availability	MTBF (days)
9020D/9020E	0.99998191	613
Rehost	0.99998922	1420

estimated availability and MTBF for each system. It is seen that the rehost system retains its edge in reliability even when software failures are taken into account with an MTBF of 1420 days vs. an MTBF of 613 days for the 9020D/9020E system. (It should be emphasized that there are some failures, e.g., failures caused by human error, that are not captured by this analysis. Therefore, the numbers in Tables ES-3 and ES-4 should be interpreted only as relative indicators of the reliability of the two systems. In other words, this analysis does not allow one to say in absolute terms what the system reliability is, but it does give a common basis for comparing the two systems. FAA data indicates that there are perhaps two or three system failures at each en route center per month.)

In summary, if the current system were replaced by the rehost system, there would be both a quantitative and a qualitative change in the reliability. Quantitatively, there would be a reduction in the number of system failures. Qualitatively, the shorter recovery times and greater predictability of the rehost system would mean that controllers would have less uncertainty about how long an interruption in service would last; this would decrease the disruption caused by short outages.

Performance. The rehost system is only of interest to the FAA if it is able to adequately handle the workload over its expected life. In order to study the question of whether the rehost system would perform adequately, this report focuses on the system response time. The idea is that the system is constantly receiving inputs such as radar data and messages from the controllers; the system's main job is to make sure that these inputs are reflected on the controller's screen in a timely manner. If performance is inadequate, then the system will fall behind the stream of inputs and the controller's screen will become out of date. Therefore, the system response time, or the time it takes an input to be reflected on the controller's screen, is the relevant way to measure performance.

Two different analyses of performance are carried out. The first is a rough global analysis which examines the degree to which modern technology surpasses that embodied in the 9020's and which concludes conservatively that the rehost system as a whole will have a service time half that of the current system. From a calculation based on queueing delays it is concluded that even if the workload were to double, the rehost system would still have a response time half that of the current response time.

The second analysis uses simple operational analysis techniques to infer system response time from the service time and utilization of individual components. This analysis proceeds in seven steps. First, characterize the typical transaction (which is a request for service such as a controller asking for information) in terms of the workload it imposes on each component. Second, assume an arrival rate for the transactions. Third, infer the percentage utilization of each component such as CPU, disk, and channel. Fourth, determine the response time for each component, using the equation that the response time equals the service time divided by one minus the utilization. (Service time is the time it would take to process a transaction if there were no congestion.) Fifth, add the response times for the CPU, disks, and channels to obtain what is termed the active server time. Sixth, determine the delays due to data base locks and non-reentrant program element locks; this is called the passive server time. Seventh, add the active and passive server times to obtain the total system response time.

The results are shown in Table ES-5. This table shows for a variety of track counts the total system response time. The first two columns show that with a track count of 110 the 9020A and the rehost system have estimated system response times of 6,735 and 128 milliseconds, respectively. The rest of the table shows that the rehost system maintains an acceptable response time for the peak track counts projected through 1995. Throughout the analysis the assumptions adopted are conservative and chosen to make sure that the response time of the rehost system is not

TABLE ES-5: PERFORMANCE OF THE 9020A AND REHOST SYSTEMS

	9020A	Rehost System Base	1980	1985	1990	1995
Track count	110	110	319	384	486	597
CPU utilization	.73	.15	.44	.52	.66	.81
CPU response time (ms)	2,263	49	74	87	123	224
Disk utilization	.38	.12	.35	.42	.53	.65
Disk response time (ms)	340	75	102	114	141	190
Channel utilization	.43	-	-	-	-	-
Channel response time (ms)	91	-	-	-	-	-
Total active server time (ms)	2,694	124	175	201	264	414
PE utilization	.60	.028	.11	.16	.26	.51
Overall response time (ms)	6,735	128	197	239	357	845

N.B. This table assumes that an IBM 3033U or an Amdahl V7 is used as the rehosting mainframe.

underestimated. Therefore, though the analysis is tentative, it does strongly suggest that the rehost system can perform more than adequately.

Technical issues. Can the NAS software be made to run successfully on the rehost machine? There are two aspects to this question.

First, the 9020's execute about fifteen special instructions that are not standard System/360 instructions and that could not be executed by the rehost machine. There are a number of different methods that could be used to deal with these instructions; the discussion indicates how these instructions could be handled by trapping and emulating the instructions, by changing the operation code, or by doing nothing since the instruction would not be executed in the rehost system.

Second, a number of features of the 9020 environment pose potential problems for the rehost system. These problems pertain to memory usage (relating to page zero, storage keys, immediate instructions, and memory size), timer usage and synchronization, program status word format, devices and channel program usage, and diagnosis and error analysis. The details of these problems and possible methods of dealing with them are discussed.

The conclusion is that the technical problems of rehosting the NAS software can be readily dealt with. While the methods sketched out might not be the best, they do at least show that suitable methods do exist.

Transition. The FAA has established the requirement that when the new machine is installed, there should be no significant interruption in the seven days a week, twenty-four hours a day provision of air traffic control services. Moreover, there must be a ninety day period in which both the old and new systems are operating so that there will be a proved back-up to the new system. In order to achieve these transition goals three problems must be dealt with.

First, there are remodeling problems since the site would have to be prepared for the new system. Second, there are technical problems since cables must be connected so that inputs can be directed to either system and so outputs can be supplied from either system; the technical problem of how the old system can take over if the new system fails must also be dealt with. Third, there are personnel problems since the training of personnel must be scheduled so that the lag between training and when the new system is installed is minimized; the training schedule, however, must prevent the center from being undermanned at any time. No detailed transition plan was developed, but an analysis of transition issues did not uncover any serious difficulties. It can be concluded that potential problems can be avoided by advance thinking and careful preparation.

Growth potential. In order to minimize the trauma of transition and to avoid the expense of repeated replacement, a system that can gradually evolve through time is desired. There are three ways that a system should be able to evolve. First, it should be able to be upgraded to include new technology as that technology becomes available. Second, it should be able to increase its capacity as the load on the system makes additional capacity necessary. These first two criteria are mainly related to hardware, and they are met since the rehost system consists of standard, off-the-shelf hardware. For example, an Amdahl V7 can be field-upgraded to a V8 over a weekend. Therefore, mainframes can be upgraded, memory can be added up to 16 megabytes, and peripherals can be replaced as desired without seriously interrupting air traffic control services. In this way the system can reflect current technology and offer increased capacity.

The third type of evolution is that the system should be able to provide additional functions as the scope of air traffic control changes. For example, the system should be able to handle the increasing levels of automation that are being introduced. This criterion is largely related to software. Since the rehost system uses the NAS software, which does not reflect modern programming practices, gradual changes to the NAS software would be difficult; in this sense, then, functional evolution of the system would not proceed smoothly. Once the rehost system is operating, however, the software could be totally rewritten, and in this way the rehost system could be put into a form that could evolve to satisfy growth in air traffic and to support fully automated air traffic control. (One outstanding issue is whether the rehost system could provide the level of reliability and availability that is needed in the long term.)

Summary. The question which the FAA is preparing to answer is: How should the procurement of a computer system to replace the 9020's proceed? This report does not attempt to address this entire question; it only looks at the pros and cons of the strategy of rehosting the NAS software on instruction-compatible machines. This report finds that the rehost system could provide improved levels of reliability and performance, and the risks due to possible technical problems and the transition appear to be acceptable. The conclusion is that the rehost system is a suitable system to adopt. This is not to say that the rehost system is the best system or that it should be adopted; this statement would require one to look not only at the rehost system but also at alternate systems, and this is beyond the scope of this report.

1. REHOSTING THE NAS SOFTWARE

1.1 Background, Purpose, and Organization of this Report

One of the missions of the Federal Aviation Administration (FAA) is to provide en route air traffic control services. To fulfill this mission the FAA has established in the continental U.S. twenty air route traffic control centers (ARTCC's), each equipped with computer systems that collect, transfer, and process the data that are used to keep current the displays and print the flight strips used by the air traffic controllers. These computer systems, along with all associated hardware and software, are known as the automation systems of the National Airspace System (NAS).

These computer systems have been in place and supporting air traffic control (ATC) for the last ten years and can be expected to provide effective support for some time to come. If air traffic increases as forecast, however, there will eventually come a time when these systems approach saturation and will not be able to keep the controllers' displays sufficiently up-to-date. Even if traffic does not increase as forecast, the age of the system, the increasing difficulty of acquiring spare parts, or the desire to have a system with greater capability means that the system will be replaced in the not too distant future.

The FAA's current plan is to fully replace the current system by 1990 with the Advanced Computer System [FAA80, p.16]. If it turns out that the current computer system cannot keep up with the growth in air traffic that takes place before 1990, then the options are to either restrict air traffic to a level that the current system can handle or to adopt an interim, short-term system that will stretch the life of the current system until full replacement can be carried out. The FAA is currently studying a number of potential interim systems.

The purpose of this report is to examine the interim system that results when the current NAS software is rehosted on an instruction-compatible machine. That is, the hardware would be replaced by modern machines that could execute the NAS software, and the software would only be changed insofar as was necessary for it to run on the new machines. This option will for brevity be referred to as "rehosting" or "instruction-compatible replacement." The potential advantages that are sometimes claimed for this option are as follows.

- The modern hardware would be fast enough to eliminate any capacity problems.

- The modern hardware would be much more reliable than the current hardware.

- Since much of the software would run on the new system without change, the time, money, and risk involved in developing new software could be largely avoided.

- Since instruction-compatible machines are available off-the-shelf, this option could be implemented quickly if it looks like capacity problems are imminent.

This report will critically examine these claimed advantages and also search for any disadvantages that might result if this option were adopted. The report is organized in the following way:

- Ch. 2: Reliability - What would be the availability of the new system compared to the current system? How often would failures occur that degraded system performance? What would be the expected duration of a system failure?

- Ch. 3: Performance - What response time can be expected from the new system compared to the current system?

- Ch. 4: Technical Issues - Can the NAS software be made to run on a modern machine with a reasonable amount of effort?

- Ch. 5: Cost - What extra cost would this option entail?

- Ch. 6: Transition - What problems would occur during the transition to the new system?

- Ch. 7: Schedule - When could the new system be in place and operating?

- Ch. 8: Growth Potential - Does the system have the capability to evolve smoothly as technology advances and as there are changes in the services that ATC provides?

The purpose of these chapters is to point out the arguments for and against the rehosting option so the FAA will have the information needed to decide whether this option should be adopted.

The remainder of this chapter describes the current computer configuration at the ARTCC's and the baseline rehost configuration. Variants on the baseline rehost configuration have been suggested, and some of them are mentioned in Sec. 1.4; these variants are not discussed in detail, however, since it is expected that they would not significantly alter the analysis.

1.2 The Current Computer Configuration

The computer system that supports the NAS at each ARTCC has two parts. First, the central computer complex (CCC) receives inputs from the radar, flight service stations, controllers, and other sources and then performs the flight data processing and the radar data processing. In other words, the CCC takes the raw information and converts it into a form that is useful to the controller. Second, the display channel takes the output from the CCC and uses it to keep each controller's plan view display (PVD) current. The CCC and display channel together, then, are responsible for taking the raw data about what is happening in the sky and providing it to the controller in a way that can be readily grasped and acted on. Figure 1-1 shows a block diagram of the current system.

The CCC at ten of the ARTCC's use an IBM 9020A system; Figure 1-2 shows the configuration of the 9020A system. Originally all the CCC's were to

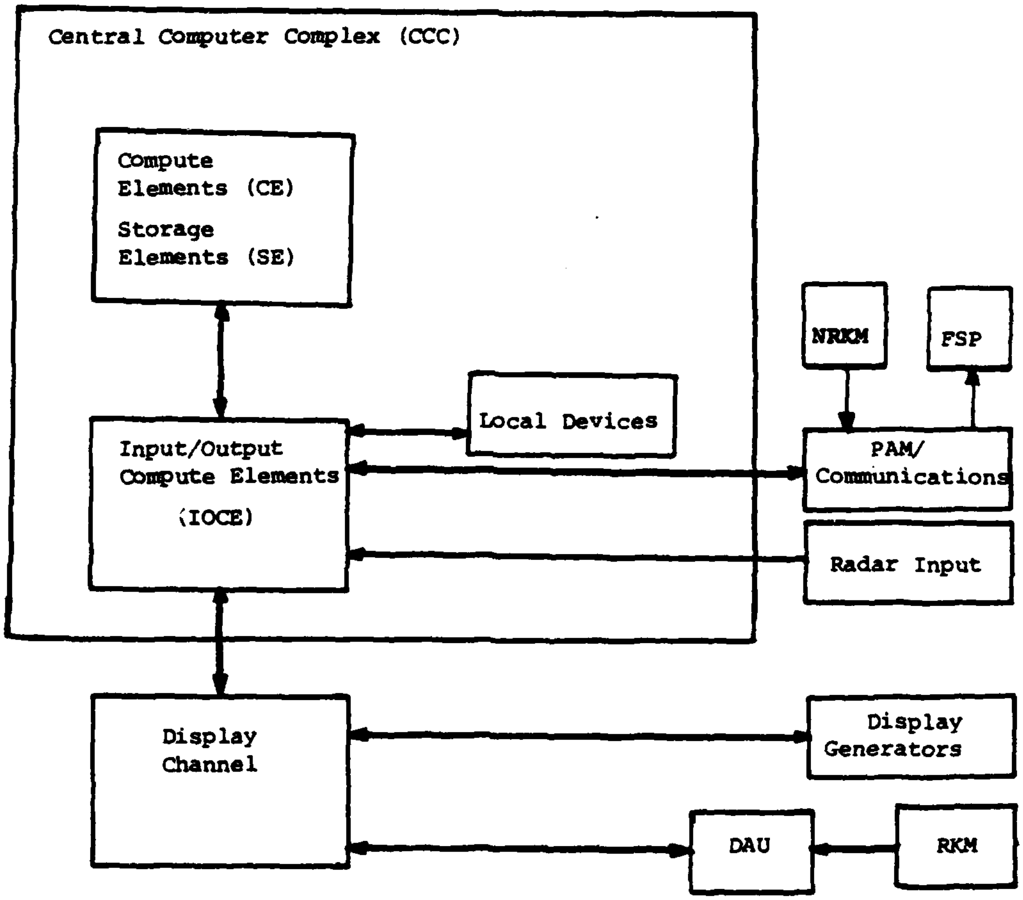

N.B. Acronyms explained in Sec. 1.3.

FIGURE 1-1: TOP-LEVEL ARTCC COMPUTER CONFIGURATION

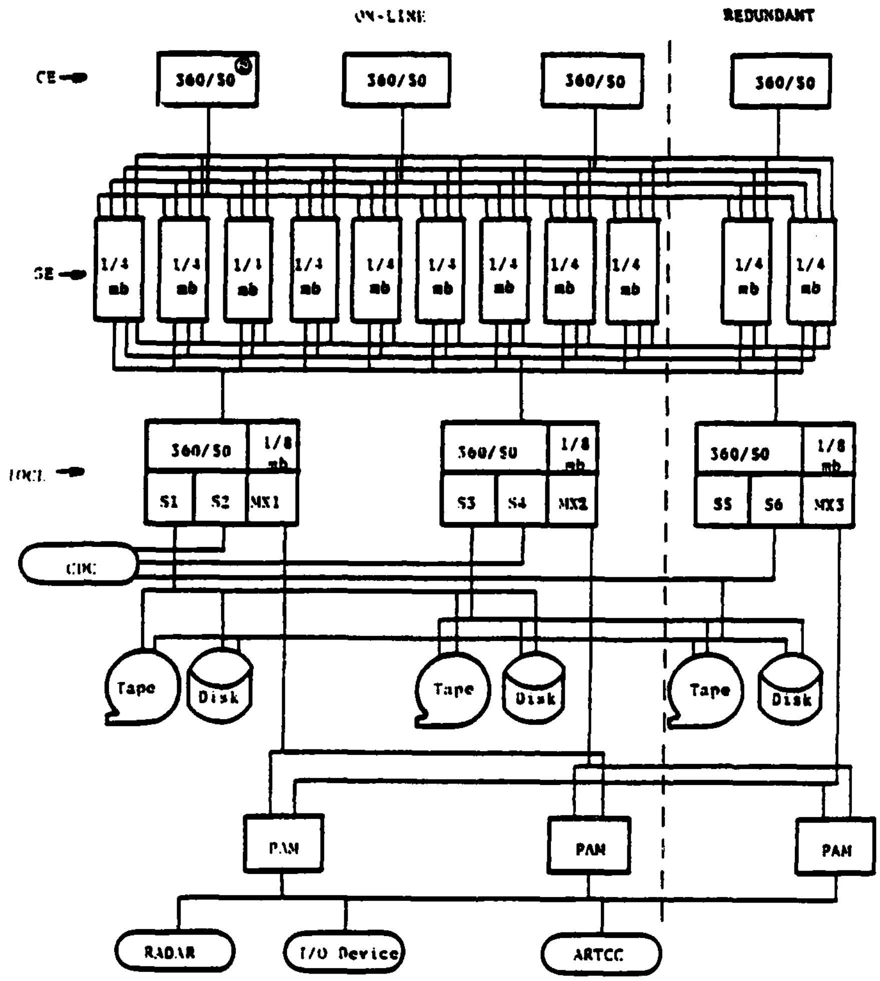

FIGURE 1-2: SIMPLIFIED 9020A CONFIGURATION DIAGRAM [CLAP79]

Si - Selector Channel
MXi - Multiplexor Channel
PAM - Peripheral Adapter Module
CDC - Display Channel

have been 9020A's, but it was feared that this would not provide enough capacity for the busier centers, so ten ARTCC's have the more powerful IBM 9020D for a CCC; Figure 1-3 shows the configuration of a 9020D. (The FAA is adding another storage element (SE) to each 9020A and 9020D; these additional SE's are not shown in Figures 1-2 and 1-3.)

The display channel at fifteen of the centers is a Raytheon 730. Since it was thought that the Raytheon 730 would not provide enough capacity for all the centers, five are equipped with a IBM 9020E in the display channel.

The term "9020's" will, somewhat inaccurately, be used throughout this report to refer to the computers in the CCC and the display channel. Table 1-1 shows which computers are in place at each ARTCC. (En route ATC is also provided at three sites outside the continental U.S. These sites are not considered in this report since they use a version of the ARTS system rather than the 9020 system. The FAA's plan is that these sites will get the same equipment as the other ARTCC sites when the Advanced Computer System is installed.)

To indicate the possible problems that the 9020's face, Table 1-2 shows the judgments made by one study as to where there are bottlenecks that potentially limit capacity. I/O bandwith, I/O device speed, and memory capacity are seen as likely bottlenecks in both the 9020A and 9020D systems; memory bandwidth and processing capacity are further bottlenecks in the 9020A system. This quick survey serves to show some of the problems that rehosting must be able to deal with. (Further information on resource usage is in [NIEL77a], [NIEL77b], and [KAND77].)

1.3 The Baseline Rehost Configuration

In order to analyze the feasibility of rehosting the NAS software in an instruction-compatible computer system, a generic, rehost computer system has been configured. This generic system is refered to as the baseline rehost system in this report and is shown in a block diagram format in Figure 1-4 and with possible components in Figure 1-5. The baseline rehost system is representative of any rehost system that would resolve the 9020's

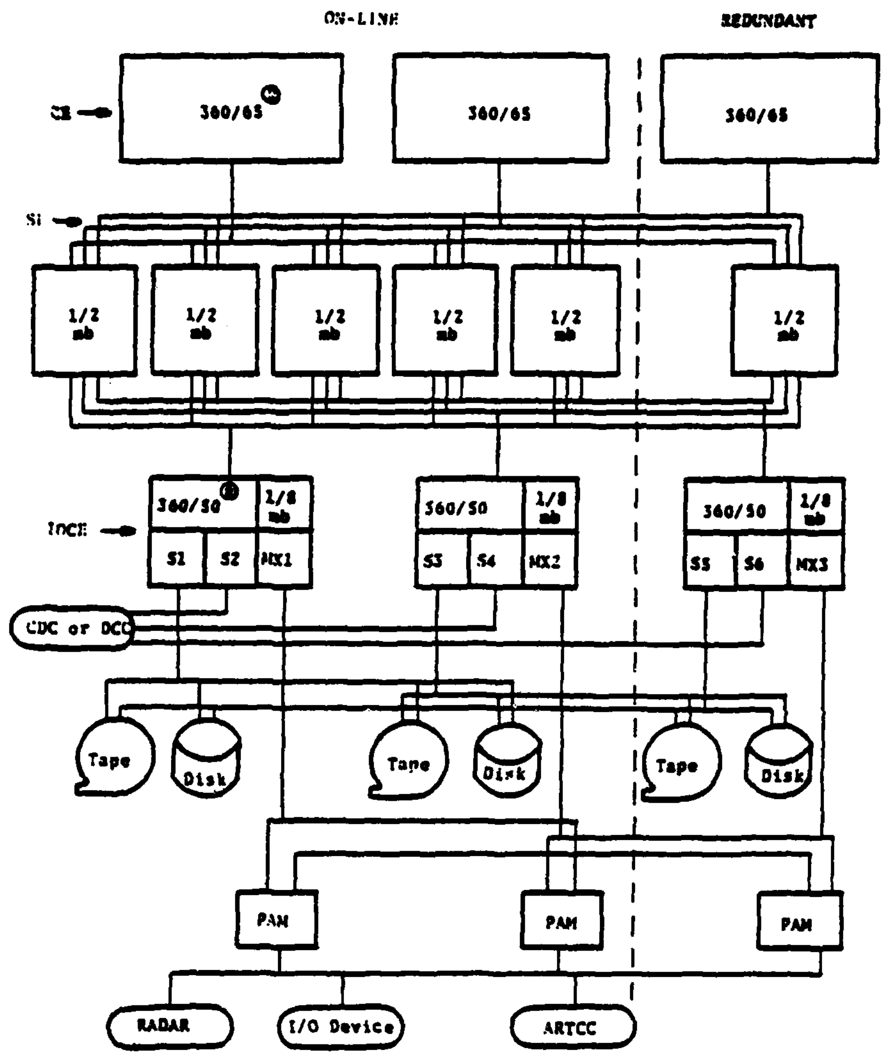

FIGURE 1-3: SIMPLIFIED 9020D CONFIGURATION DIAGRAM [CLAP79]

Si - Selector Channel
MXi - Multiplexor Channel
PAM - Peripheral Adapter Module
CDC/DCC - Display Channel

TABLE 1-1: COMPUTER SYSTEM CONFIGURATIONS FOR THE ARTCC'S

Center	CCC	Display	Sectors
Albuquerque	IBM 9020A	Ray 730	34
Atlanta	IBM 9020D	Ray 730	41
Boston	IBM 9020A	Ray 730	32
Chicago	IBM 9020D	IBM 9020E	43
Cleveland	IBM 9020D	IBM 9020E	47
Denver	IBM 9020A	Ray 730	34
Fort Worth	IBM 9020D	IBM 9020E	39
Houston	IBM 9020A	Ray 730	41
Indianapolis	IBM 9020D	Ray 730	34
Jacksonville	IBM 9020D	Ray 730	37
Kansas City	IBM 9020D	Ray 730	36
Los Angeles	IBM 9020D	Ray 730	37
Memphis	IBM 9020A	Ray 730	36
Miami	IBM 9020A	Ray 730	28
Minneapolis	IBM 9020A	Ray 730	34
New York City	IBM 9020D	IBM 9020E	39
Oakland	IBM 9020A	Ray 730	39
Salt Lake City	IBM 9020A	Ray 730	21
Seattle	IBM 9020A	Ray 730	22
Washington DC	IBM 9020D	IBM 9020E	36

TABLE 1-2: THE CRITICAL 9020 RESOURCES

	Is this resource a bottleneck?	
Resource	9020A	9020D
I/O Bandwidth	Yes	Yes
I/O Device Speed	Yes	Yes
Memory Capacity	Yes	Yes
Memory Bandwidth	Yes	No
Processing Capacity	Yes	No

Source: [CLAP79, p. C-20]

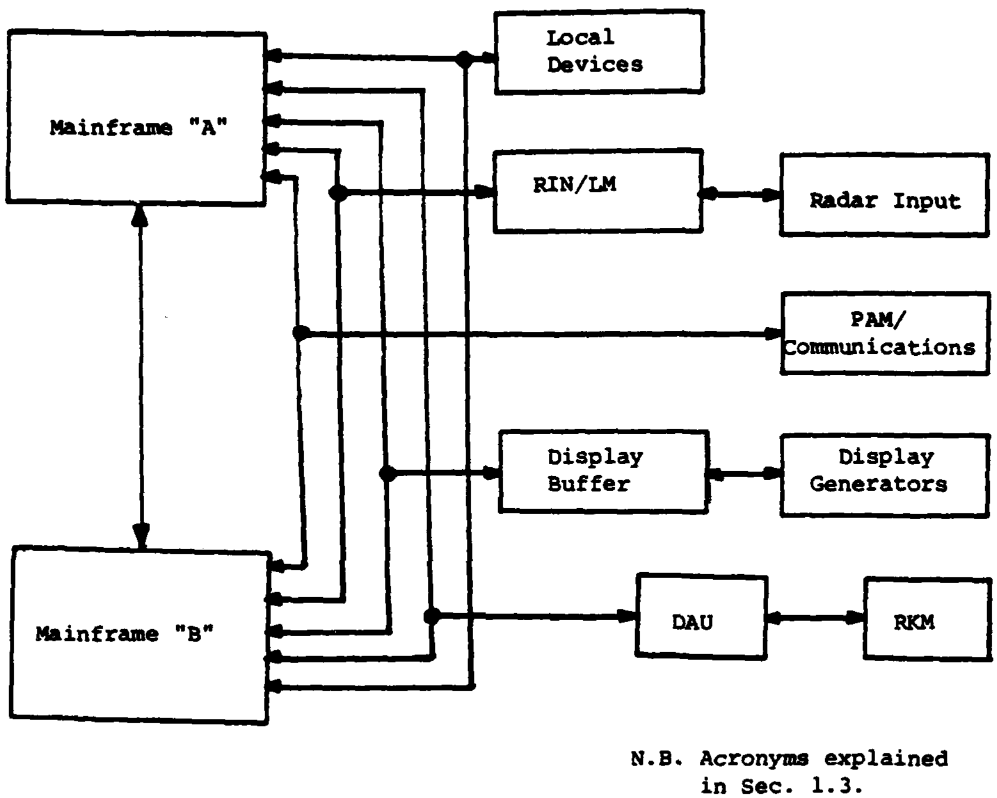

FIGURE 1-4: BASELINE REHOST CONFIGURATION

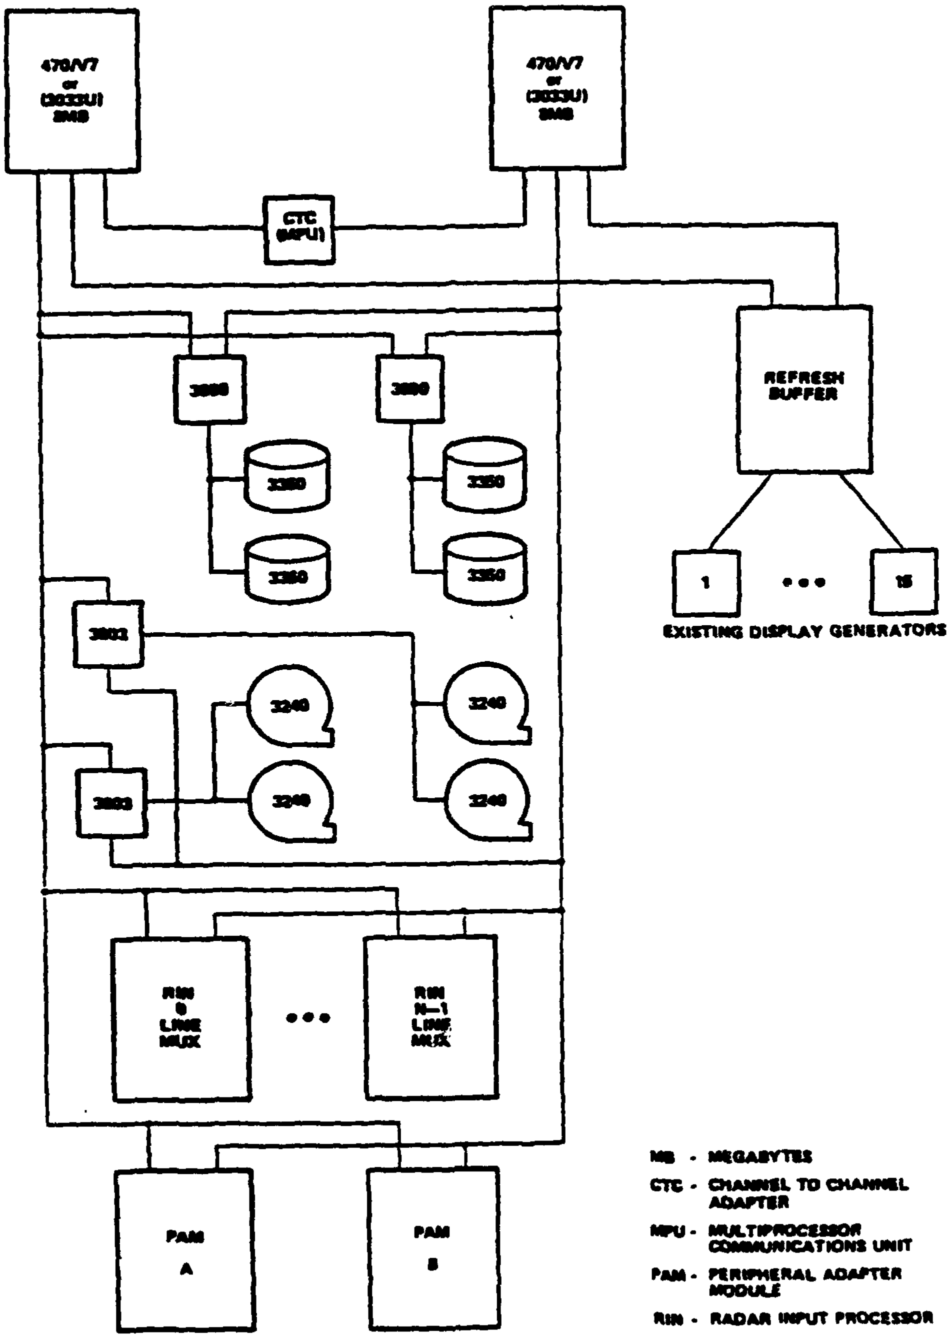

Source: TASC

FIGURE 1-5: DETAILED BASELINE REHOST CONFIGURATION bottlenecks and satisfy the operational constraints of the ARTCC facilities. The following is a list of constraints and precepts that any rehost computer system must satisfy:

- The rehost system will replace both the CCC (9020D and 9020A systems) and the display channels (Raytheon 730 and 9020E systems).
- There will be two mainframe computers in the rehost configuration, and they will be used in a duplex mode of operation.
- Each mainframe will be capable of supporting all of the current CCC and display channel processes.
- The two mainframes will be connected via a channel-to-channel adapter.
- The interface to the controller suites will be at the display generators and the keyboard control units.
- The interface to the radar data input circuits will be at the circuit terminations.
- The interface to the various communications circuits will be at the peripheral adapter module (PAM).
- The interface to the radar keyboard multiplexer (RKM) will be at the data adapter unit (DAU).
- All the devices local to the processor (for example, disk, tape, line printers, operator console, and terminals) will be replaced.
- The controller suites, the flight strip printers (FSP), and the non-radar keyboard multiplexors (NRKM) will be retained without change.
- DARC will provide an independent radar data channel capability for backup purposes.

- All devices will be connected to both mainframes and have two paths to each mainframe.

- Each unique device type will be represented by at least two devices.

These general properties are not enough to define the rehost system completely. Therefore, to complete the characterization of the baseline rehost system, this report adopts the following assumptions.

- Each mainframe will have eight megabytes of main memory with a growth potential to sixteen megabytes.

- Each mainframe will have twelve channels that can operate either in multiplexor or block multiplexor mode. These channels are divided into six pairs, where one channel in each pair is redundant. Figure 1-4 shows how these six pairs are connected.

- The normal mode of operation would be for one mainframe to support all of the processes while the other is maintained in a "hot standby" status.

- A radar input line multiplexor (RIN/LM) will serve each radar data input circuit and block valid radar data into records for processing by the mainframe. RIN/LM is described in App. D.

- A display buffer will be located between the mainframe and each display generator to provide the necessary display file memory for the display generators and avoid memory contention problems in the mainframe. The display buffer is described in App. D.

The baseline rehost configuration is a representative configuration rather than the only or best configuration. Permutations to this configuration are described in the next subsection. General operational procedures for the rehost system are described after the discussion of permutations to the configuration.

The cost, transition, and schedule analysis assumes that the rehost system is deployed at the twenty ARTCC's plus the FAA Technical Center (two systems) and the FAA Aeronautical Center, for a total of twenty-three systems. App. C discusses the cost saving if there is partial replacement.

1.4 Permutations to the Baseline Rehost Configuration

The baseline rehost configuration provides a basis for discussing the advantages and disadvantages of the concept of rehosting. There are, however, many variants on the baseline that should be mentioned. The leading variants and some of their features are as follows.

- Replace only the CCC component of the current computer system and retain the current display channels. There are several problems with a CCC only replacement.

  + Actual experience with DARC and the 9020E has led to the conclusion that the CCC/display channel interface is more complex than the display channel/display generator interface.

  + A display channel outage would result in a system outage since the current display channels are redundant on a component basis but not on a unit basis.

  + The queueing delays for the Raytheon 730 display channel [NIEL77a] would remain unresolved.

  + The ongoing maintenance costs for the current display channels would exceed the incremental cost of replacing the display channels as part of the CCC replacement.

- Retain the current 2314 disks to avoid any embedded channel program problems that could arise with new disks and their associated device support routines. Several benefits of current technology disks would not be available to the rehost system:

+ shorter access and latency times than those for 2314 disks,

+ higher data transfer rates than those for 2314 disks,

+ better reliability characteristics than those for 2314 disks,

+ larger storage capacity than that for 2314 disks.

- Replace the PAM's and DAU's with newer technology control units and line controllers. The engineering costs for redeveloping all of the necessary line controllers and the problems associated with the physical transition to the replacement PAM's and DAU's must be weighed against current maintenance costs for these units and the need for flexibility.

- Consider supporting RIN in the mainframe. RIN is currently supported with an open loop channel program in the IOCE. This sort of channel program would not be viable in a mainframe. However, the other alternative of interrupt driven radar data input for RIN would destroy the performance of the mainframe.

A programmatic variant to the baseline rehost configuration is to deploy the rehost system at only the overloaded ARTCC's. This partial deployment would reduce the rehost hardware procurement costs but would require logistics and maintenance support for two very different systems. In addition, a partial deployment of the rehost system would make more difficult an orderly evolution of the ATC functions.

1.5 Operation of the Baseline Rehost Computer System

In order to manage the resources of the mainframe and to support local data processing activities, it is expected that a virtual machine environment will be provided in the rehost mainframes. VM/370 represents a viable virtual machine monitor for this application. Another option is use the kernel of VM/370 as a basis for developing a virtual machine monitor unique to the needs of rehosting the NAS software and providing a virtual environment for supporting the local data processing activities as well as the evolving ATC functions. Whatever monitor is used for the mainframe, it must support:

- the application component of the NAS software without revision and with a minimum of monitor overhead;

- the local utility programs for data analysis, report generation, adaptation assemblies, and system generation.

There are two possible modes for operating the rehosted NAS software since the baseline rehost configuration will have two mainframes to satisfy availability requirements and each mainframe will have sufficient processor capacity to support both the CCC processes and the display channel processes. One mode of operation is to designate one mainframe as the "active" processor and the other as the "standby" processor with automated support for transferring active status in the event of a failure. The other mode of operation is to assign the CCC processes to one processor and the display channel processes to the other processor. In this split mode of operation, a failure of one processor would result in the transfer of all processes to the operational processor. The analysis is based on the first mode of operation because it will simplify the backup procedures.

REFERENCES

[CLAP79] David Clapp, Joseph Hagopian, and Ronald Rutledge, Analysis of Expandability and Modifiability of Computer Configuration Concepts for ATC, Volume I, report no. FAA-EM-79-22, November 1979.

[FAA80] "The FAA Plans and Programs for the Future Airport and Traffic Control System: Presentations at the Air Traffic Control Computer Replacement Industry Briefing," prepared by the FAA, December 2, 1980.

[KAND77] William Kandler, Gary Nielsen and John Squires, Response Time Analysis Study Fort Worth ARTCC Measurement, Interim Report, Logicon report no. R4940-107, August 1977.

[NIEL77a] Gary Nielsen, William Kandler and John Squires, Response Time Analysis Study Memphis ARTCC Measurement, Interim Report, Logicon report no. R4940-107, July 1977.

[NIEL77b] Gary Nielsen, William Kandler and John Squires, Response Time Analysis Study Indianapolis ARTCC Measurement, Interim Report, Logicon report no. R4940-107, June 1977.

2. RELIABILITY

2.1 Purpose and Organization of this Chapter

The reliability of an air traffic control computer system is an important criterion in deciding whether it should be procured since system outages can cause delay and cancellation of flights as well as a decreased level of safety and an increased workload on the controllers. The purpose of this chapter is to discuss the reliability of the rehost system compared to that of the 9020 system.

The two main questions of interest that this chapter focuses on are: How often does a system failure occur? How long will it last? In considering these questions it is important to remember that the failure of a single unit of hardware does not cause a system failure. This is because both the 9020 and rehost systems have redundant hardware and the capability to automatically reconfigure the system so that the interruption in the operation of the computer system is only a matter of seconds when an individual unit fails. For example, the 9020D has three compute elements (CE's); under normal operation two are active and one is redundant. If one of the active CE's fails, then the 9020D automatically reconfigures so that the redundant CE is made active; this process takes perhaps 25-30 seconds. This means that even if one component fails, there is no system failure, i.e., no significant interruption in service, because of the redundancy built into the system. In this example there would only be a system failure if, before the first failed CE were repaired or replaced, a second CE failed. This shows how redundancy can lessen but not eliminate system failures.

This chapter is organized as follows. Sec. 2.2 shows how the system availability, system mean time between failure (MTBF), and average duration of a system failure can be estimated from information on the MTBF of individual units, the mean time to repair (MTTR) individual units, and the configuration of the system. The system availability, system MTBF, and the expected duration of the system outage are estimated for the 9020D/9020E system and the rehost system. A sensitivity analysis is carried out that shows how these estimates vary if alternate assumptions are used. This analysis only considers hardware and does not discuss software reliability. Sec. 2.3 looks at hardware reliability from another angle by assuming that no repairs are made; system MTBF's are estimated, and a sensitivity analysis is carried out.

Sec. 2.4 gives a qualitative discussion of the software reliability that could be expected from the rehost system. Sec. 2.5, in a tentative, quantitative analysis, then goes on to make numerical assumptions about software reliability, and estimates the availability, MTBF, and expected duration of an outage for the system as a whole that takes into account both hardware and software.

Finally, Sec. 2.6 discusses the failures that result from miscellaneous problems such as errors by human operators and technicians.

A cautionary note should be sounded about the theoretical nature of the results reported in this chapter. <u>If</u> the assumptions made about component MTBF's and repair time are correct, <u>then</u> the results in this chapter are valid. The data that is available to check these assumptions, however, is incomplete; therefore, there is doubt about the accuracy of the assumptions. Moreover, it would have been desirable to validate the model by checking the results against the measured MTBF's and availability of the 9020's, but the available data were too incomplete to allow this. For these reasons the reader should reserve judgment on the accuracy of this chapter's results. Like the miles per gallon figures featured in automobile advertisements, these results are to be used only for purposes of comparison.

Some terminology is needed. Availability is the amount of time a system is working divided by the sum of the time the system is working and the time the system is not working. Equivalently, availability is the probability that the system is working at a randomly chosen point in time. The terms operating, working, up, and not failed are used synonymously. It is assumed that a unit or a system is either failed or not failed; no intermediate stage of partial failure is considered. The terms component, unit, element, and device are used synonymously to refer to a single storage element (SE), compute element (CE), input/output compute element (IOCE), tape control unit (TCU), or storage control unit (SCU, i.e., disk). All of the like units are referred to as a subsystem, e.g., the three CE's in a 9020D are the CE subsystem.

2.2 Hardware Reliability

2.2.1 Introduction

This section, which considers hardware only, presents estimates of the system availability, system MTBF, and the expected duration of a system outage for the rehost system and for a system with a 9020D in the CCC and a 9020E in the display channel. The exposition proceeds in four steps.

- Develop a theoretical model which expresses system availability, system MTBF, and the expected duration of a system outage as a function of the unit MTBF's and unit mean time to repair.
- Determine the MTBF's and MTTR's for each component.
- Substitute these MTBF's and MTTR's into the model to obtain estimates of system availability, system MTBF, and the expected duration of a system outage.
- Carry out a sensitivity analysis to determine how the results vary if alternate assumptions are used.

Each of these steps is now discussed.

2.2.2 The Model of System Availability and System MTBF

App. A presents a detailed derivation of the equations and methods used to estimate system availability and MTBF. This subsection describes this analysis. It has three main steps.

First, determine the configuration of each system to be modeled. The two systems modeled are the rehost system and a system with a 9020D in the

CCC and a 9020E in the display channel. (Information on the Raytheon 730 was not sufficient to allow it to be modeled.) Consider the rehost system. Define a mainframe to consist of a CPU, a memory, and six pairs of channels. A mainframe is working if the following three conditions all hold:

- the CPU is working;
- the memory is working;
- at least 1 channel in each pair is working.

The rehost system is working if the following three conditions all hold:

- at least 1 of the 2 mainframes is working;
- at least 1 of the 2 TCU's is working;
- at least 1 of the 2 SCU's is working.

For the 9020D, the system is working if the following five conditions all hold:

- at least 2 of the 3 CE's are working;
- at least 5 of the 6 SE's are working;
- at least 2 of the 3 IOCE's are working;
- at least 2 of the 3 TCU's are working;
- at least 2 of the 3 SCU's are working.

It is assumed that for this analysis the 9020D and 9020E are equivalent.

Second, derive the equations that express system reliability. For a single unit the equation for its availability A_u is

$$A_u = \frac{MTBF_u}{MTBF_u + MTTR_u} .$$

From the availability of a single unit, an equation is derived that states the availability of a subsystem, e.g., the probability that at least 2 of the 3 9020D CE's are working at a randomly chosen point in time. From the availability of the subsystems, the availability of the complete system is derived, i.e., the probability that the system is working at a randomly chosen point in time.

Third, the system MTBF is derived. If the equation above is interpreted to apply to the system instead of just a unit and is solved for MTBF, one gets

$$MTBF_s = \frac{A_s}{1-A_s} MTTR_s.$$

Since system availability A_s was calculated in the second step, system MTBF can be estimated once the system MTTR is calculated. The system MTTR, which is the same thing as the expected duration of a system outage, is estimated with a special calculation. This completes the outline of how information about unit MTBF's, unit MTTR's, and the configuration can be combined to estimate the system availability, MTBF, and MTTR. (It is important to note that the variables in the first equation refer to a single unit while those in the second refer to the entire system.)

In summary, the main assumptions used in this derivation are:

- all failures are probabilistically independent;
- the MTBF for each unit is finite;
- all repair times are independent and exponentially distributed;
- when an active unit fails and is replaced by a redundant unit, the reconfiguration is instantaneous;
- a system failure only occurs under the conditions spelled out above;
- no more than one of these conditions is violated at any one time.

There are three ways in which these assumptions are not exact. First, repair times are not independent since the time it takes to repair one unit is affected by whether, when it fails, there are any other units being repaired. Second, reconfiguration after a unit failure is not instantaneous; in the 9020D it takes perhaps 25-30 seconds. Third, it is possible that two of the conditions could be simultaneously violated; this is, however, an extremely unlikely event. None of these three simplifying assumptions has a significant effect on the results.

2.2.3 MTBF and MTTR Data

In order to estimate system availability and system MTBF, the model just described requires data on the MTBF and MTTR for each component. For the rehost system, MTBF's are taken from [RUTL81]. The original source of this data is Reliability Research, Inc., which gathers data on the actual reliability of equipment operated by a variety of users. The MTBF's refer to machines comparable to those that would probably be used in a rehost system. The MTBF's used are shown in Table 2-1. The column labeled "Best Estimate" are taken from [RUTL81]. Low and high estimates, which are needed for the sensitivity analysis, are also shown; the low and high estimates are, respectively, half and twice the best estimate.

For the 9020D system, all MTBF's except that for the SCU were taken from [MOSS75, p. 25] which reports failure data on the 9020D at the FAA Technical Center over the one-year period starting October 1, 1971. For the CE, SE, and IOCE this study computes the lower and upper bounds of the 95 percent confidence interval; these bounds are displayed in Table 2-1 along with the best estimate. No interval was computed for the TCU since no failures were observed; Table 2-1 uses for the low and high estimates, respectively, half and twice the number given in [MOSS75]. The MTBF for the SCU is 25,358 hours in [MOSS75]; this is not a representative figure since disks were just being introduced during the period of observation and were not fully utilized. Therefore, the MTBF for the SCU is taken from [RUTL81]. [RUTL81] is not a good source for the other 9020D MTBF's because, while figures for System/360 components are given, they are based on extremely small samples.

In sum, the MTBF's used in this study are shown in Table 2-1. These are the best figures that could be obtained, but it should be stressed that there are real doubts as to the accuracy of these figures. For example, this table shows that a 9020D CE has a higher MTBF than a modern CPU; this is counter to the generally accepted opinion that modern technology is much more reliable than System/360 technology. A sensitivity analysis using alternate MTBF's is carried out to try to minimize this problem, but a sensitivity analysis is not a substitute for good data. Therefore, because non-comparable and perhaps inaccurate data are used, one should reserve judgment on the accuracy of the results reported.

TABLE 2-1: MTBF's USED IN THIS STUDY

	MTBF (hours)		
9020D Components	Low	Best Estimate	High
CE	1,391	2,301	4,116
SE (1/2 MB)	2,089	3,173	5,052
IOCE	1,750	3,161	6,354
TCU	4,241	8,482	16,964
SCU	350	700	1,400
Rehost Components			
CPU	678	1,356	2,712
Memory (2MB)*	1,147	2,293	4,586
Channel	658	1,316	2,632
Tape	500	1,000	2,000
Disk	3,582	7,163	14,326

* Since each rehost mainframe has 8 megabytes of main memory, the MTBF's used in the calculations are one-fourth of the figures shown in this table.

Source: [MOSS75, p. 25] and [RUTL81]

The time to repair each component is assumed to be exponentially distributed with a mean of one hour. This assumption was chosen after talking to Airway Facilities Service personnel and after examining the data in [MOSS75, p. 25]. Since the data on repair times is scanty, this assumption should be treated as tentative. It is also assumed that the repair time for any unit is independent of whether any other units are failed. In effect, this assumption means that failed units need not queue up waiting for a repairman.

2.2.4 Estimates of System Availability and System MTBF

When the best estimates of component MTBF's are substituted into the model's equations, the estimates of system availability and system MTBF shown in Table 2-2 result. The mean time between system failures for the rehost system is 2905 days, which is about 2 1/2 times the MTBF of 1226 days for the 9020D/9020E system. Therefore, in this primary calculation the rehost system is substantially more reliable than the 9020D/9020E system. For both systems the expected duration of a system outage is a half hour.

The main reason why the rehost system has a higher MTBF lies in its configuration; the two rehost mainframes are configured in parallel whereas the CCC and the display channel in a 9020 system are configured in series. That is, if one rehost mainframe fails, this does not cause a system failure. But if either the 9020D or the 9020E fails, this does cause a system failure. A numerical example will bring out the importance of this consideration. Suppose that the probability that a particular mainframe is working is 0.5. Then the probability that at least one mainframe is working is $1 - (0.5)^2 = 0.75$. Now consider the 9020D/E system, and suppose that the probability of the 9020D working is 0.5, and that the probability of the 9020E working is also 0.5. Both must work to prevent a system failure, and

TABLE 2-2: SYSTEM AVAILABILITY AND SYSTEM MTBF: PRIMARY CALCULATION

System	System Availability	System MTBF
9020D/9020E	0.99998301	1226 days
Rehost	0.99999283	2905 days

the probability of both working is then $(0.5)^2 = 0.25$. This example indicates how the parallel configuration of the rehost system increases system availability.

Another factor that should increase the relative reliability of the rehost system is that it embodies modern technology, which is much more reliable than the technology embodied in the 9020's. For example, when an instruction fails to execute in a 9020, there is a machine check. In contrast, when an instruction fails to execute in the rehost system, the machine check is held pending and the instruction is retried; there is a machine check only if the instruction fails to execute twice. As Table 2-1 shows, this advantage of the rehost system is not fully reflected in data used in this study. Because of this apparent flaw in the data, the reported results probably understate the reliability of the rehost system compared to the 9020D/9020E system.

2.2.5 Sensitivity Analysis

To show how the results given in Subsec. 2.2.4 are affected by variations in the data, the results of a sensitivity analysis will now be presented. Two types of variations are considered. First, to recognize the uncertainty in the unit MTBF's, the calculation is carried out not only for the baseline MTBF's but also for the low and high unit MTBF's shown in Table 2-1. Second, to recognize the uncertainty in the unit MTTR, the calculation is repeated using not only the baseline MTTR of 1 hour but also the alternate values of 1/2 and 2 hours.

The results for the 9020D/9020E system are shown in Table 2-3 under 39 sets of assumptions. The first line labeled "Baseline" used the baseline unit MTBF's from Table 2-1. The second line labeled "High CE" used the high CE MTBF from Table 2-1; the baseline MTBF's are used for the remaining units. The rest of the cases, with the exception of the last two, similarly use the baseline MTBF's for all but one component; the table shows for which component an alternate MTBF is used and whether the alternate MTBF is the high or low value. The next to last line uses the high MTBF's for all the components; the last line used the low MTBF's for all the components.

TABLE 2-3: 9020D/9020E SYSTEM AVAILABILITY AND MTBF: SENSITIVITY ANALYSIS

Component MTBF's	MTTR = 1/2		MTTR = 1		MTTR = 2	
	Avail.	MTBF	Avail.	MTBF	Avail.	MTBF
Baseline	0.99999575	2449	0.99998301	1226	0.99993226	615
High CE	0.99999594	2566	0.99998379	1285	0.99993536	645
Low CE	0.99999525	2195	0.99998105	1099	0.99992442	551
High SE	0.99999619	2739	0.99998481	1372	0.99993945	688
Low SE	0.99999477	1993	0.99997913	998	0.99991676	501
High IOCE	0.99999586	2516	0.99998346	1260	0.99993406	632
Low IOCE	0.99999541	2268	0.99998165	1136	0.99992684	569
High TCU	0.99999576	2458	0.99998307	1231	0.99993251	617
Low TCU	0.99999568	2413	0.99998276	1208	0.99993126	606
High SCU	0.99999804	5306	0.99999215	2655	0.99996866	1330
Low SCU	0.99998660	778	0.99994660	390	0.99978791	196
High All	0.99999881	8760	0.99999525	4383	0.99998102	2195
Low All	0.99998474	682	0.99993915	342	0.99975818	172

N.B. MTBF is in days.

For each case the results are shown for the three assumptions for the MTTR of 1/2, 1, and 2 hours. For example, this table shows that if the high value for the IOCE MTBF of 6,354 hours is used, and the baseline MTBF's are used for the remaining components, then the system MTBF is 2516, 1260, or 632 days depending on whether the assumed MTTR is 1/2, 1, or 2 hours. Table 2-4 shows the results for the rehost system under 39 sets of assumptions; this table is read the same way as Table 2-3. Five conclusions can be drawn from the sensitivity analysis.

First, the system MTBF's are very close to being inversely proportional to the unit MTTR. This holds for both the 9020D/9020E system and the rehost system. For example, for the 9020D/9020E system, with the baseline MTBF's and an MTTR of 1/2 hour, the system MTBF is 2449 days; when the MTTR is doubled to 1 hour, the MTBF is nearly halved to 1226 days; when the MTTR is again doubled to 2 hours, the MTBF again is nearly halved to 615 days.

Second, results for the 9020D/9020E system are very sensitive to the SCU MTBF but relatively insensitive to the other unit MTBF's. For example, consider the case where the MTTR is 1 hour. Under the baseline unit MTBF's, the 9020D/9020E system MTBF is 1226 days. If the SCU MTBF is then raised from its baseline value of 700 hours to 1400 hours, the system MTBF rises to 2655 days, which is almost as high as the rehost system MTBF of 2905 days under the baseline assumptions. If the SCU MTBF is instead lowered from 700 to 350 hours, then the 9020D/9020E system MTBF falls to 390 days. No such wide swings occur when the other unit MTBF's are varied; in the other cases in which a single unit MTBF is changed, the 9020D/9020E system MTBF falls into the interval from 998 to 1372 days.

Third, the rehost system MTBF is relatively insensitive to changes in the MTBF's of the channels, TCU's, and SCU's. The explanation is that these units have such high MTBF's that they almost never cause a system failure; this remains true even after the unit MTBF's have been halved or doubled.

Fourth, the rehost system MTBF is very sensitive to the unit MTBF's of the CPU and the memory, especially the latter. For example, if the memory MTBF is halved from 2,293 to 1,147 hours, the rehost system MTBF falls to

TABLE 2-4: REHOST SYSTEM AVAILABILITY AND MTBF: SENSITIVITY ANALYSIS

Component MTBF's	MTTR = 1/2		MTTR = 1		MTTR = 2	
	Avail.	MTBF	Avail.	MTBF	Avail.	MTBF
Baseline	0.99999821	5807	0.99999283	2905	0.99997136	1455
High CPU	0.99999863	7599	0.99999452	3801	0.99997809	1902
Low CPU	0.99999716	3665	0.99998865	1835	0.99995470	920
High Memory	0.99999909	11536	0.99999639	5766	0.99998554	2880
Low Memory	0.99999529	2214	0.99998123	1110	0.99992533	558
High Channel	0.99999821	5813	0.99999284	2911	0.99997146	1460
Low Channel	0.99999819	5787	0.99999278	2885	0.99997095	1434
High TCU	0.99999839	6484	0.99999358	3244	0.99997434	1624
Low TCU	0.99999746	4098	0.99998984	2051	0.99995947	1028
High SCU	0.99999821	5819	0.99999284	2911	0.99997141	1458
Low SCU	0.99999819	5760	0.99999277	2882	0.99997112	1443
High All	0.99999955	23222	0.99999821	11615	0.99999283	5811
Low All	0.99999283	1454	0.99997137	728	0.99988577	365

N.B. MTBF is in days.

1110 days, which is less than the 9020D/9020E system MTBF of 1226 days under the baseline assumptions.

Fifth, to summarize, if the same MTTR is assumed for both systems, then the rehost system's highter MTBF in the baseline case is maintained throughout most of the cases examined. The conclusion is that the rehost system's higher MTBF is not highly sensitive to the assumptions used here. It should be pointed out, however, that the rehost system's lead in system MTBF can be reduced and even lost if one picks and chooses from among the cases so that the rehost system is put in the worst light and the 9020D/9020E system in the best light.

2.3 Hardware Reliability: The No Repairs Case

One of the uncertainties in the analysis of the previous section is doubt over what the repair time would be. The sensitivity analysis that assumed various mean repair times is one way of dealing with this doubt. Another way of dealing with it is to assume that repairs are not made; that is, the system runs until enough unit failures have accumulated to cause a system failure. This approach is taken in this section. This approach is flawed because the assumption that repairs are not made is incorrect, but it does give one a way of comparing the different systems on a common basis which is uncontaminated by a possibly inaccurate assumption about repair time.

The reliability of a system R(t) is defined to be the probability that after t hours of operation there has not been a system failure. App. B derives equations that allow the reliability function R(t) to be derived once the MTBF's of the individual units are known. The main assumption is that the failure time for each unit is exponentially distributed. The reliability function, once it is obtained, can then be used to estimate the system MTBF, i.e., the number of hours the system is expected to operate before a system failure halts operation. (The system MTBF thus estimated is not exact; it is approximated by a procedure explained in App. B.)

Using the best estimates of the unit MTBF's in Table 2-1, it is found that the rehost system has an MTBF of 400 hours, and the 9020D/9020E system has an MTBF of 300 hours. Therefore, in this analysis the rehost system retains its lead in reliability.

A sensitivity analysis is also carried out for this no repair approach, and the results are shown in Tables 2-5 and 2-6. The tables show the system MTBF in hours for a variety of cases. The baseline case uses the best estimates of Table 2-1. Each succeeding case indicates the individual component for which the MTBF is varied, and high or low tells which of the alternate MTBF's from Table 2-1 is used. For example, in Table 2-5 the "High CE" case means that the high CE MTBF of 4116 hours from Table 2-1 is used; all other MTBF's are the baseline MTBF's. That is, except for the cases marked "All" at the bottom of the tables, only one unit's MTBF is changed for each calculation. For the cases marked "All," the MTBF's of every component are changed for the calculation. Examination of these tables shows that, for the most part, the rehost system maintains its edge in reliability; it is, however, possible to pick cases in which the 9020D/E system has a higher system MTBF.

2.4 Software Reliability

The reliability of software is an important aspect of any computer system since the hardware and software are combined in a serial manner to support every application. That is, failure in either the hardware or the software will result in a system failure. This means that perfect hardware alone cannot overcome software defects and conversely. Before proceeding to a quantitative analysis in Sec. 2.5, this section will give a qualitative discussion of what the reliability of the rehost system software is expected to be compared to the current software.

In the rehost system, there are three major software components:

- the NAS application software,
- the NAS monitor, and
- the virtual machine monitor (VMM).

TABLE 2-5: 9020D/9020E SYSTEM MTBF'S WITHOUT REPAIRS

Case	9020D/9020E MTBF(hours)
Baseline	300
High CE	310
Low CE	280
High SE	320
Low SE	270
High IOCE	310
Low IOCE	290
High TCU	300
Low TCU	300
High SCU	420
Low SCU	180
High All	560
Low All	160

TABLE 2-6: REHOST SYSTEM MTBF'S WITHOUT REPAIRS

Case	Rehost MTBF(hours)
Baseline	400
High CPU	440
Low CPU	350
High Memory	490
Low Memory	300
High Channel	490
Low Channel	280
High TCU	420
Low TCU	360
High SCU	400
Low SCU	400
Low All	200

The NAS application software should have reliability characteristics in the rehost system that are equivalent to those for the current NAS application software since this software will not be changed in the rehosting process. Some additional application software failures are expected during the testing phase as problems in the interfaces to the new NAS monitor components and the VMM are identified. These interface problems are expected to be resolved before operational use. It is important to note that rehosting the NAS application software will preserve its current reliability characteristics and cannot improve them. However, rehosting the NAS software will allow changes in the usage of that software which would result in improvements in its reliability. For example, the large memory of the rehost system will allow all program elements and tables to be memory-resident and will avoid problems with program element and table buffering. This will improve reliability since swapping into and out of main memory is currently a significant source of software failures.

The NAS monitor will be modified as part of the rehosting process to accommodate changes in the hardware and the system configuration. These changes will degrade the initial reliability characteristics of the NAS monitor in the rehost system. After some period of operational usage, the reliability characteristics of the modified NAS monitor can be expected to return to the current level of reliability.

The VMM represents a new software component; unless the VMM were perfect with respect to reliability, the VMM would result in some degradation of the overall software reliability. The possible range of failures for a VMM is indicated by the current commercial experience with a large virtual memory operating system. That is, this system has 1 to 3 failures per month [RUTL81, p. 2-11].

The net effect of rehosting on the overall reliability of the NAS software is that the reliability is expected to be about equal to the current reliability after some period of operational usage since the benefits of memory-resident NAS application software will be largely offset by defects in the VMM.

2.5 System Reliability

Sec. 2.2 produced quantitative estimates of the hardware availability and MTBF for the rehost system and the 9020D/9020E system. Since, however, system reliability depends both on hardware and on software, it is desirable to extend the analysis to include not only hardware but also software. The problem is that the information needed to include software in the analysis is not available. Nevertheless, because of the interest in the reliability of the total system and not just the hardware, some ballpark assumptions about software will now be made so that quantitative estimates can be made of system availability, system MTBF, and the expected duration of a system outage. It must be stressed that these assumptions made about software do not have a solid foundation; they are made here for illustrative purposes. Six assumptions about software are used.

- The number of system failures caused by software is equal to the number caused by 9020D/9020E hardware (based loosely on FAA operational experience).

- When the NAS software fails, with probability 0.9 the failure is transient and the system outage during the dynamic recovery is exponentially distributed with a mean of 30 seconds. With probability 0.1 the system must be restarted, and the resulting system outage is exponentially distributed with a mean of 15 minutes.

- The VMM fails at the rate of twice per month.

- When the VMM in a mainframe fails, the system outage (while processing is transferred to the other mainframe) is exponentially distributed with a mean of 10 seconds. (This assumes that critical data are saved every five seconds and all software is loaded and ready to run in the back-up system.)

- When the VMM fails in a mainframe, the time that the mainframe is down while the VMM is restarted is 10 minutes.

- The NAS software has the same MTBF in the rehost system as in the 9020D/9020E system.

The assumptions made about hardware are that the unit MTBF's are the best estimates in Table 2-1 and that the repair times are independently, exponentially distributed with a mean of 1 hour.

This treatment assumes that the NAS software is combined serially with the hardware; a failure in either causes a system failure. The VMM, however, is treated like a component in a mainframe, e.g., just like the CPU. When the VMM fails, processing is transferred to the other mainframe, and the first mainframe is in failure mode until the VMM is restarted.

The availability, MTBF, and expected duration of a system failure that are implied for each system are calculated in App. A and are shown in Table 2-7. The system MTBF for the rehost system of 1420 days is more than twice that of the 9020D/9020E system of 613 days. While this is admittedly a rough calculation, it is sufficient to refute the claim that the improved reliability of the rehost hardware would be cancelled out by the unimproved NAS software.

2.6 Non-Standard System Failures

In addition to the hardware and software failures discussed so far, there are also miscellaneous failures that are grouped together under the heading of non-standard failures. These failures arise from:

TABLE 2-7: AVAILABILITY, MTBF, AND EXPECTED DURATION OF A SYSTEM OUTAGE

System	Availability	MTBF (days)	Expected Duration of a System Outage (minutes)
9020D/9020E	0.99998191	613	16.0
Rehost	0.99998922	1420	22.0

Source: App. A

- operator errors,
- technician repair errors, and
- exogenous events (e.g., earthquakes).

It seems likely that exogenous events would have a similar effect both on the current systems and on the rehost system.

Good data on operator and technician errors is apparently not available, but discussions with FAA personnel indicate that human error accounts for a large percentage of failures. Since the rehost system would be much more reliable, we would expect a significant reduction in these failures. That is, with fewer failures, there would be fewer problems requiring operator intervention or repairs, and there would, therefore, be fewer opportunities for these types of failures. It is not possible to go beyond this qualitative statement because of lack of data and understanding of these non-standard failures.

2.7 Summary

This chapter has shown that in terms of reliability the rehost system has both advantages and disadvantages when compared to the current system. The main advantage of the rehost system is that, because it uses duplex processors and modern technology, it has significantly greater hardware reliability. While the analysis has not been verified empirically, the increased reliability of the rehost system is so pronounced that it seems unlikely that this result could be reversed by any changes to the analysis. The main reliability disadvantage of rehosting is that it would require an additional software component, the virtual machine monitor, and this would present a continuing software reliability problem. Other software reliability problems would result because of the changes to the NAS monitor, but it is expected that these problems would decline in importance after an initial shakedown period. In short, the rehosting system would show an increase in hardware reliability and be nearly equal in software reliability.

The analysis of Sec. 2.5 suggests (but does not prove) that the rehost system will offer a significantly greater reliability even after it is recognized that the NAS code will still be used and that, in addition, a virtual machine monitor will be used and will be a source of failures. It should be emphasized that this rough analysis does not provide any definitive answers, but it does provide a systematic way of thinking about the question of how rehosting would affect system reliability.

REFERENCES

[MOSS75] Arthur R. Moss, "Reliability Analysis of the 9020A and 9020D Central Computer Complex at the National Aviation Facilities Experimental Center," prepared for the FAA, report no. FAA-RD-75-175, November 1975.

[RUTL81] Ronald M. Rutledge, A Survey of Hardware Failures for Large Scale IBM 9020 Instruction-Compatible Computers, Transportation Systems Center, March 1981.

3. PERFORMANCE

3.1 Purpose and Organization of this Chapter

The purpose of this chapter is to estimate the response time that the rehost system could provide both in an absolute sense and also compared to the 9020's. In order to estimate the response time two separate analyses are carried out.

The first analysis, which is described in Sec. 3.2, treats the entire system as a single, undifferentiated server. The extent to which technological progress has increased hardware speed and capacity is discussed, and a decrease in service time provided by the system is determined. From this the system response time is estimated. This rough analysis shows that rehosting does make sense from a performance point of view; a more detailed analysis, therefore, is justified.

The second analysis uses the technique of operational analysis to look at the response time of each individual component as a function of its service time and utilization. The response times of individual components are then added to obtain the system response time. Sec. 3.3 explains the principles of this technique, and Sec. 3.4 applies it to estimate response times.

Throughout this chapter when the data is ambiguous or unsatisfactory, conservative assumptions are used. Therefore, if anything, the response time of the rehost system would be better than the conservative estimates made here.

3.2 A Global Performance Analysis

3.2.1 Introduction

The 9020 systems are modified versions of the IBM 360 series of computers, which was designed in the early sixties and introduced into the market in the mid-sixties. The 9020 systems, therefore, generally use hardware technology and programming methodologies that were developed in the 1960-1965 timeframe. During the intervening two decades, technology has progressed rapidly, and this trend will continue in the foreseeable future. For instance, semiconducter chip technology has advanced by five orders of magnitude (10^5) in the 1965-1980 timeframe.

If a decision to rehost the NAS software is made in the near future, it would take several months in calling for bids, evaluating them and making a final choice between the alternative systems. In our analysis, we therefore include systems that would become available during 1981 and 1982.

In the succeeding paragraphs, we consider various system elements separately.

3.2.2 Central Processing Units

The processing capacity of the 9020A CE (7201-1) and the 9020D or 9020E CE (7201-2) have been identified [WHAI, p. 3-2] as 286 KOPS (kilo-operations per second) and 1,000 KOPS, respectively. These values have been used for the performance calculations and, while they are different from the KOPS values [LIAS80, p. 104] used elsewhere in this report, these differences will not significantly affect the results of the performance analysis.

IBM models now on the market that are upward-compatible with System/360 span the spectrum from 2,300 KOPS through 22,200 KOPS. Thus the CPU speed accelerator factor is 1.6 to 15.2 times the performance of a 9020A system, which consists of 3 9020A processors and 2 IOCE processors; it is 0.9 to 8.6 times that of a 9020D system, which consists of 2 9020D processors and 2 IOCE processors. Taking an average, the speed increase factor can be up to 12 times the performance of a 9020 system.

3.2.3 Memory Units

The 9020A and the 9020D systems are today equipped with on-line memory units aggregating 2.25 megabytes and 2.50 megabytes, respectively. The upper limit on currently available systems is generally either 16 megabytes or 32 megabytes, with newer systems of up to 64 megabytes expected in the next one to two years. Thus, memory capacity of either a 9020A or a 9020D can be enhanced by up to 25 times by using newer memory hardware. On the speed front, however, the performance improvement is not so dramatic. The 9020A storage element has a cycle time of 2.5 microseconds for an access width of 4 bytes, and the 9020D storage element has a cycle time of 0.8 microseconds for 8 bytes. Today's systems exhibit memory cycle times of 0.3-0.4 microseconds per 8 bytes, and thus the speed increase factor is between 2 and 8. However, almost all new systems offer a fast cache memory whereas the 9020 did not implement a cache; these caches have capacities of up to 64 kilobytes, and their cycle time is between 50-100 nanoseconds. As the cache-hit ratio nears 100%, the effective memory cycle time becomes equal to the cache cycle time. This would represent a speed-up factor of up to 15 over a 9020D memory system. On the whole, and preferring to err on the conservative side, we expect memory speed increase to be between 2 and 10.

3.2.4 Disk Units

The existing IBM 2314 disk units have a maximum capacity of 30 megabytes and a maximum data transfer rate of 625 kilobytes a second. The disk service time of 34 milliseconds is composed primarily of the access time and the seek time (both mechanical functions) and, by comparison, a small data transfer time, especially for small data blocks. For example, a 2K data block has a total transfer time of 34 milliseconds, of which 3 milliseconds is transfer time. Today, 300 megabyte disks are common, and 1000 megabyte disks have recently become commercially available. Thus the capacity increase factor for disks is between 10 and 30. The disk transfer times have improved from 625 kilobytes/second to about 2000 kilobytes/second, a factor of 3. The improvement factor for access times and seek times is between 1 and 2 only. However, since the primary memory would be much larger than that of the existing system, a substantial number of programs, and possibly some flight plan data, could under rehosting reside in main memory, thus greatly reducing the number of disk accesses. With this revised system design, we estimate that the overall speed increase in disk service time would at least be between 2 and 3.

3.2.5 Tape Units

The technology trend for tapes is similar to that for disks. The capacity increase factor will be between 5 and 10, and the overall speed increase factor between 2 and 3. Since the disk units will now have much higher capacities, it is possible to store archival data on disks rather than on tapes as is presently done. This will improve the speed increase factor.

3.2.6. Other CCC Equipment

Aside from the devices considered in the foregoing paragraphs, the CCC consists of input devices, peripheral adapter modules (PAM's), and channels. Computer systems today offer 4-16 channels per CPU, each with a capacity of 2-10 megabytes per second; besides, the facility of block-multiplexor mode, in addition to the traditional selector and multiplexor modes, will mitigate any channel bottlenecking. The 9020A, as well as the 9020D, have two PAM units each. The load on all these units is still below their capacity limits, and taking into account the technological improvements, it is unlikely that there would be any problem in this area from the performance viewpoint.

3.2.7 Display Equipment

In the baseline rehost configuration, the display channels (Raytheon 730 and 9020E) would be replaced while the display generators and controller suites would be retained. The processing capacity necessary to support the display channel functions is very low. That is, current estimates of the display processor utilization range from 1% to 12%. The effects of the display channel workload on the rehost system will be minimal.

3.2.8 Computational Workloads

A report prepared by the Transportation Systems Center [CLAP79] indicates that the air traffic volume has been increasing at the rate of 4.4% per annum, a doubling every 15 years. This implies that if no system enhancements are made during the intervening period, the computational workload in 1990 would be roughly twice that which existed in 1975, and the 1995 workload would be twice the 1980 workload. The rehost system design could be modified to benefit from the larger capacities of primary memory, disks and tapes; this would enable some near-term enhancements in NAS system capabilities to be introduced without degradation of system performance. As such, it would still be appropriate to assume a doubling of workload every 15 years as a ballpark figure (assuming no change in level of automated ATC services or demand.)

3.2.9 Rough Calculations

Assuming that the arrival rates and the service times are exponentially distributed, the response time of a given server, or device, can be calculated using the formula:

$$\text{Response Time} = \frac{\text{Service Time}}{1-\text{ (Service Time) x (Arrival Rate)}} . \qquad (1)$$

For a fixed arrival rate, this formula shows that if the service time doubles, the response time will more than double; likewise, if the service time falls by 50%, the reduction in response time would exceed 50%.

The analysis of subsections 3.2.2 through 3.2.7 is summarized in Table 3-1. This table shows that the speed improvement factor is between 2 and 50 depending on the nature of the device, or that the individual service times will reduce by somewhere between 50% and 98% of the respective existing times. In the most conservative case, assume that the reduction is 50% for all devices, electronic, mechanical, or whatever.

Over a 15 year timeframe, the arrival rate of transactions is expected to double, hence the product of service time and arrival rate (defined as utilization) will remain constant. Thus, in equation (1), the denominator will remain constant, and the numerator will be halved, hence the time interval from a request for service to the completion of the service (response time) will be reduced by 50%. The reduction factor means the response time of the rehosted system in 1995 will be one-half the response

TABLE 3-1: TECHNOLOGY TRENDS

Device	Speed Increase Factor *	Capacity Increase Factor **
CPU	Up to 12	Same Instruction Set
Memory	2-10	Up to 25
Disk	2-3	10-30
Tape	2-3	5-10

* Compared to devices used in the 9020 system

** Compared to maximum capacity of a 9020A or 9020D system

time of the existing system today. Since we have assumed very conservative technology factors throughout, and have deliberately preferred to err on the safe side while calculating acceleration factors, it can be concluded that rehosting would maintain an acceptable response time under the expected air traffic levels.

The above analysis also indicates that it is not really necessary to use the most advanced and most capable current technology computer system to achieve rehosting. We now focus on a typical-sized computer that should be adequate for rehosting the NAS software.

3.3 An Operational Analysis of Performance: Principles

3.3.1 Overview

This subsection provides a scenario for the analysis of the performance of the rehost system in terms of its resource utilizations and response times. The prime source of the data used is a collection of relevant reports on the performance analysis of the existing systems [KAND77], [NIEL77a], [WHAI81]. A detailed analysis requires more data on the workload and the future system characterization than are available at this time due to a different orientation of the performance recording of the existing systems and the lack of a benchmark on the future rehost system. An attempt is made to use available data to project the performance of the rehost system through a set of qualitative analyses. Moreover, as pointed out in Sec. 3.2, the utilization of the present display channels (CDC or DCC) is relatively low, and as such its contribution to the total computational workload is minimal; the impact of transferring these functions to the rehosted system would also be marginal. Therefore, this performance analysis will concentrate on the workload of the CCC.

Throughout this section the operational analysis technique [DENN78] is employed. The simplest form of this technique embodies the following equation:

$$R = S/(1-u),$$

where R is the response time of a certain type of workload, S is the service time, and u is the utilization of the resource in question. The resource may be an active server (e.g., CPU, channel, devices) or a passive server (e.g., data base, program elements). The operational analysis technique relaxes the restrictions on the distributions of the arrival rates and the service times. The only assumption used is that the flow of transections through the system is balanced, which is satisfied in the system being evaluated.

The method used in evaluating the performance of the rehost system is summarized below.

(1) Identify the characterizations of a <u>typical</u> rehost system.

(2) Identify a base scenario workload: the characterizations of typical transactions and their arrival rates under a specified air traffic load on the present 9020A system.

(3) Conduct an analysis to obtain the utilization rates and response times of the rehost system under the work load of the base scenario identified in (2) above. This analysis will be done by integrating the utilization rates and response times calculated for each resource type in the rehost system.

(4) Compare the numbers obtained in (3) with the performance figures of the present system to derive improvement ratios.

(5) Extrapolate the impact of the increased air loads on the arrival rates of the transactions and perform a sensitivity analysis of the rehost system performance.

During the process of analysis, several assumptions are made where data is lacking. These are described as the analysis proceeds.

3.3.2 Characterization of a Typical Rehost System

The hardware configuration of the rehost system has been briefly described in Sec. 1.3. The CPU speed of the rehost machine is assumed to be 5900 KOPS, such as found in current generation mainframes such as the IBM 3033U or Amdahl V7 [LIAS80, p.104]. However, due to the sensitivity of this speed factor to the cache-hit ratio of the CPU when running the NAS programs, we assume a 5% degradation of speed performance to arrive at an effective speed of 5605 KOPS. The memory size of the rehost system will be at least 8 megabytes, with a growth potential of up to 16 megabytes. As the size of the present NAS software is estimated to be around 4.1 megabytes, it is expected that in the rehost system all buffered program elements (PE's) and buffered flight plan data are to be memory-resident, thus eliminating the need for swapping. The choice of the memory size of the rehost system should aim at elimination of swapping. Future increase in the size of the NAS software due to functional enhancements or increase in the size of the data bases should be taken into account in deciding the size of the memory of the rehost system.

The channels in the rehost system will have the block multiplexing capability and a much higher transfer rate (e.g., 2.6-6 Megabyte/Second). The disk and tape units will also be replaced by modern-technology counterparts.

The rehost system will run VM/370 to ease the environmental changes for the NAS monitor and the application software. Current commercial experience indicates that VM processing results in about 25% CPU overhead, which reduces the effective speed of the rehost CPU to 4203 KOPS.

Since software rehost minimizes modification to current NAS software, non-reentrant PE's and queueing delays due to PE or database lockups will still exist.

Other input and output devices such as non-radar-keyboard and flight strip printers will be retained. With the exception of flight strip printers and FDEP's these devices are not considered highly utilized and will not significantly contribute to the response times. Therefore our analysis will concentrate on the CPU, channel, disk and program utilizations.

3.3.3 Characterization of a Typical Transaction

A transaction is characterized by its resource service time and arrival rate. Normally the transactions processed by a system are grouped into a small number of classes; transactions within each class consume similar amounts of resources and have other similar properties (e.g., priorities). However, the present system recording in the 9020's does not provide resource consumption on a per transaction basis. Therefore, for the purpose of a preliminary performance analysis, we have aggregated all resource utilizations and distributed them among all input transactions (including radar and timer messages) to derive the resource consumption of a "typical" transaction. To do so, we make use of the aggregate data provided by [WHAI81].

CPU and Disk Time

[WHAI81] provides the following data for the Houston 9020A site:

Transaction arrival rate = 12,960 per hour

CPU utilization = 73% (or 2.2 out of 3 CE's)

Disk utilization = 38% per disk

Track count = 110

From the above data, we derive the following:

CPU time per transaction = 611 ms

Disk time per transaction = 211 ms

The CPU time per transaction is derived as follows:

(1) Total CPU time = (3,600,000 ms/hour) x (number of CE's busy)
(2) CPU time per trans. = Total CPU time/arrival rate,

and the disk time per transaction is derived as follows:

(1) Total disk time = (3,600,000 ms/hour) x (disk utilization x 2)
(2) Disk time per trans. = Total disk time/arrival rate.

3.4 An Operational Analysis of Performance: Results

3.4.1 Scenario Analysis

This subsection presents an analysis of the performance of the rehost system when running the base scenario workload described in the previous subsection.

CPU performance. A comparison between the CPU speeds of the present systems and the CPU speed of the rehost system is made to derive the speed ratios shown in Table 3-2.

The adjusted rehost CPU speed has been derived by taking into account the best estimates for VM overhead (25%) and the cache-hit ratio degradation (5%)

Employing the basic operational analysis equation, $R = S/(1-u)$, the figures in Table 3-3 are obtained which characterize the CPU response time per transaction in the present and the rehost systems.

Disk time. The present 2314 units have an access time comparable to that of the 3330's, the replacement disks. However, the disk utilization will be dramatically reduced in the rehost system due to the elimination of buffered program elements and flight plan data bases. The Logicon studies provide the information on the disk activities shown in Table 3-4.

Based on this observation, it is assumed that the disk utilization will be reduced by 68.6% for the 9020A systems. These are translated into comparable reductions in the disk times per transaction in these systems and the characterizations for the disk activities in the rehost system are shown in Table 3-5.

Channel time. The utilizations of the two selector channels in the present system are directly related to the disk and tape activities. While the channel time is not considered significant and therefore not fully analyzed in the Wilson-Hill study, the following predictor equation was given in the Logicon report [KAND77, p. 3-22]:

Channel Utilization % = Disk u. x (25 ms/access time) + SAR% + REMON%

= Disk u. x (25 ms/access time) + (.0732 x Active) + 6.35,

TABLE 3-2: PROCESSOR SPEED COMPARISONS

	System		
	9020A	Rehost	Adjusted Rehost
CPU Speed (KOPS) per CE	286	5900	4203
Speed Improvement Ratio*	14.7	.71	1

* normalized by the adjusted rehost speed.

TABLE 3-3: CPU RESPONSE TIME PER TRANSACTION FOR THE 9020A AND THE REHOST SYSTEM

	System	
	9020A	Rehost
CPU time (ms)/ trans.	611	41.6
CPU utilization (per CE)	0.73	0.15
CPU response time (ms)	2263	49

TABLE 3-4: INFORMATION ABOUT DISK ACTIVITY FOR THE 9020A AT MEMPHIS

	Percentage of Total Disk Activities		
Track Count	Buffered PE	Buffered Flight Plans	Total
124	50.0	18.6	68.6

Source: [NIEL77a]

TABLE 3-5: DISK RESPONSE TIME FOR THE 9020A AND THE REHOST SYSTEM

	System	
	9020A	9020A Rehost
Disk time (ms)/trans.	211	66.25
Disk utilization (per disk)	38%	12%
Disk response time (ms)	340	75

where Active is the active flight account, and, from the same report, is found to be approximately 1.5 times the track count. Also it is estimated by Logicon that the average disk access time is 38 ms [KAND77, p. 3-11]. Based on the above discussion and analysis in the previous paragraphs, the figures shown in Table 3-6 are derived.

In the rehost system the reduction of disk usage combined with the introduction of the block multiplexor channels is expected to dramatically reduce the channel utilization. Furthermore, the channels to be used in the rehost system will have a transfer rate up to 2.6 megabytes per second, approximately 6.5 times that of the present system. It is therefore concluded that the channel wait time in the rehost system as a percentage of the total response time will be negligible, and it is ignored in our response time analysis.

Passive servers. By adding up the CPU and the disk response times and the channel wait time presented above, one obtains the expected response time per "typical" transaction without regard to output device I/O delays and delays due to data base locks and non-reentrant PE locks. Because the output devices are in general not to be replaced in the rehost system and their service times are not included in the response time definitions as specified in NAS-MD-318, they will not be considered in this performance analysis. However, the PE and data base locks are potential contributors to the response times in both the present systems and the rehost systems.

Judging by Wilson-Hill's experience in performance modeling, the PE lock delay is expected to be substantial, while the data locks do not contribute significantly to the overall response time. A pessimistic assumption is made for the purpose of analysis that the PE's as a passive server for our "typical" transaction have an upper bound of 60% utilization. This means that this passive server has an average service time equivalent to the aggregated average active server times per transaction, and is, under current load, 60% utilized. The purpose of this worst case assumption is to predict how the rehost system will perform under this adversity. That is, the overall response times, taking into consideration the PE and data base locks, are derived from the total active server times using the PE utilization and are shown in Table 3-7.

TABLE 3-6: CHANNEL WAIT TIME FOR THE 9020A*

Channel time (ms)/trans.	120
Channel utilization (per channel)	43.3%
Channel wait time (ms)	90.5

*"Wait time" is defined to be the response time minus service time.

TABLE 3-7: OVERALL RESPONSE TIMES

	System		
	9020A	Rehost 9020A	Improvement Ratio
CPU response time/trans. (sec)	2.263	0.049	46.2
Disk response time/trans. (sec)	0.340	0.075	4.5
Channel wait time/trans. (sec)	0.091	-	
Total active server time (sec)	2.694	0.124	27.2
PE utilization	60%	2.8%	27.2
Response time/trans.* (sec)	6.735	0.128	52.6

* including wait for PE locks

Note that this passive server's service time is proportional to the total active server time, and therefore under the same air load will be very sensitive to the technology used by the active servers.

3.4.2 Sensitivity Analysis

This subsection presents an analysis of performance of the rehost system under varying air traffic loads. The purpose is to project the rehost system performance into the 1990 timeframe based on the traffic load predictions. It was evident from the Logicon and the Wilson-Hill studies that the 9020A systems are already occasionally failing to provide adequate services under today's air traffic load. Their performance in the 1990's is not analyzed here.

The basic assumption underlying the present analysis is that the arrival rate of "typical transactions" is largely proportional to the track count handled by the center [PRES81, p. 3-1].

As all ARTCC's are required to handle the peak traffic load with adequate performance, the projected peak track counts presented in [APO81] are used as a basis for projection. Note that these numbers represent traffic load in the busiest center in the country; the average centers will be handling peak track counts much lower than these. Note also that in our base scenario presented in the previous subsection, the track count was 110 while the 9020A CPU utilization was 73%. Comparing this data with that reported in the 1977 Logicon report on the Memphis, 9020A site, which cites a 66% CPU utilization with a track count of 124, it seems that the CPU workload per transaction in our base scenario is on the high side. Therefore, the projected CPU utilization is expected to be on the high side. The results are shown in Table 3-8. These calculations show that the performance of the rehost system will remain satisfactory through the middle of the 1990's. Note that as the load increases, the actual performance will be increasingly more sensitive to the validity of the parameters and assumptions used in the base scenario analysis. Since the base scenario uses conservative estimates and assumptions, the actual performance of the rehost system in the 1990's is likely to be far better than that presented. However, since the CPU is the bottleneck in the analysis for Table 3-8, another analysis based on the assumption that the rehost system would have a 10,000 KOPS CPU is presented in Table 3-9.

The CPU could be upgraded in many ways for the rehost system to accommodate unexpected growth in air traffic or uncertainty in the parameters for the analysis. For example, one candidate CPU for the rehost system, the Amdahl 470/V7, can be field-upgraded to a model V8 and raise the gross processing capacity of the rehost system from 5,950 KOPS to 6,375 KOPS [LIAS80, p. 104].

TABLE 3-8: PERFORMANCE PROJECTION OF THE REHOST SYSTEM[1]

		Rehost System[2]				
	9020A	Base	1980	1985	1990	1995
Track count	110	110	319[3]	384[3]	486[3]	597[3]
CPU utilization	.73	.15	.44	.52	.66	.81
CPU response time (ms)	2,263	49	74	87	123	224
Disk utilization	.38	.12	.35	.42	.53	.65
Disk response time (ms)	340	75	102	114	141	190
Total active server time (ms)	2,694	124	175	201	264	414
PE utilization	.60	.028	.11	.16	.26	.51
Overall response time (ms)	6,735	128	197	239	357	845

1 This prediction is conservative in that it is designed to be the worst case prediction for rehosting.

2 A 5900 KOPS CPU is assumed, e.g., an IBM 3033U or an Amdahl V7.

3 [APO81] These figures are the forecasts of the peak track count, which by definition is the largest track count sustained over a seven minute period. In every case the peak is at the Chicago ARTCC; peaks at most of the other ARTCC's are considerably smaller.

TABLE 3-9: PERFORMANCE PREDICTION OF A REHOST SYSTEM (WITH A 10,000 KOPS CPU)[1,2]

	Base Scenario	1980	1985	1990	1995
Track count	110	319[3]	384[3]	486[3]	597[3]
CPU utilization	.09	.26	.31	.40	.48
CPU response time (ms)	27	33	36	41	48
Disk utilization	.12	.35	.42	.53	.65
Disk response time (ms)	75	102	114	141	190
Total active server time (ms)	102	135	150	182	238
PE utilization	.023	.088	.12	.18	.29
Overall response time (ms)	104	148	170	222	335

1 This prediction is conservative in that it is designed to be the worst case prediction for rehosting.

2 The IBM 3081 and Amdahl 5860 will provide at least 10,000 KOPS.

3 [APO81]

REFERENCES

[APO81] "FAA Forecast of Air Route Traffic Control Center IFR Aircraft Handled and Instantaneous Airborne Counts: FY 1981-FY2011," Office of Aviation Policy and Plans, FAA, June 1981.

[CLAP79] David Clapp, Joseph Hagopian, and Ronald Rutledge, Analysis of Expandability and Modifiability of Computer Configuration Concepts for ATC, Volume I, Transportation Systems Center, Report No. FAA-EM-79-22, November 1979.

[DENN78] Peter Denning and Jeff Buzen, "Operational Analysis of Queueing Network Models," ACM Computing Surveys, Vol. 10, no. 3, September 1978, pp. 225-262.

[KAND77] William Kandler, Gary Nielsen, and John Squires, Response Time Analysis Study Fort Worth ARTCC Measurement, Interim report, Logicon, Inc., Report No. R4940-107, Aubust 1977.

[LIAS80] Edward J. Lias, "Tracking the Elusive KOPS", Datamation, November 1980, pp. 99-105.

[PRES81] Jacque Press, "Computer Utilization at Several Enroute Air Traffic Control Centers (A302.10 System)," Report No. ARD-140-8-81, June 1981.

[NIEL77a] Gary Nielsen, William Kandler, and John Squires, Response Time Analysis Study Memphis ARTCC Measurement, Interim report, Logicon, Inc., Report No. R4940-107, July 1977.

[WHAI81] Wilson-Hill Associates, Inc., "Interim Findings Report, Performance and Analysis of FAA's En Route Air Traffic Control Computer System," prepared for the FAA, April 1981.

4. TECHNICAL ISSUES

4.1 Introduction

Chapters 2 and 3 have examined the levels of reliability and performance that could be achieved by rehosting. These chapters have implicitly assumed that the current NAS software can indeed be made to run successfully on the replacement machine. The purpose of this chapter is to identify the problems that might keep the rehosted software from running and to indicate how these problems might be dealt with. Sec. 4.2 discusses the special instructions executed by the 9020 and Sec. 4.3 discusses special features of the 9020 environment.

4.2 Special Instructions

4.2.1 Introduction

The 9020 computers are capable of executing all of the standard IBM System/360 instructions plus several special instructions [IBM73]. These special instructions are shown in Table 4-1 along with the number of times each occurs in the NAS CCC code. These special instructions are essential for the operation of the NAS software in a multi-processor and multi-processing enviroment. However, as indicated in Table 4-1, their static usage is quite low (less than 0.1% of all the instructions in the NAS software). In addition, the usage is confined to about 10% of all modules [FAAT81] and most of these modules support startup, startover and diagnostic functions.

4.2.2 NAS Application Software

The usage of 9020 special instructions in the NAS application software has been investigated at the FAA Technical Center [FAAT81] as part of their effort to demonstrate that the flight data processing (FDP) subsystem of the NAS software could be rehosted on an IBM 4341 computer. The results of this investigation are that only three special instructions are directly used in the NAS application software; they are:

TABLE 4-1: 9020 SPECIAL INSTRUCTIONS AND THEIR USAGE

Instruction	Mnemonic	Op Code	NAS Usage	Comments
Set configuration	SCON	01	12	
Delay	DLY	0B	62	
Load identity	LI	0C	16 (est)	
Set address translation	SATR	0D	5	
Insert address translation	IATR	0E	15	Conflict with MVCL
Load data address	LDA	99	1	IOCE only
Start IOCEp	SIOP	9A	16	
Set PCI	SPCI	9B	0	
Store PS base register	SPSB	A0	5	
Load PS base register	LPSB	A1	18	
Move word	MVW	D8	428	
Convert and sort symbol	CSS	02	?	9020E
Convert weatherline	CVWL	03	?	9020E
Repack symbol	RPSB	0F	?	9020E, Conflict with CLCL
Load chain	LC	52	?	9020E

Source: [IBM73] and [IBM75]

* Delay,
* Load identity,
* Move word.

The remaining special instructions support supervisor needs in a multi-processor system and occur only in the NAS monitor.

The support for these special instructions can be provided in many ways in an instruction-compatible computer; one approach based on the FDP demonstration [FAAT81] is:

* Delay: Trap the operation code and suspend the program element (PE) for the specified delay interval. Delay is used for synchronizing modules (not needed in a uniprocessor environment) and for accommodating communication circuit transients.

* Load identity: Trap the operation code and return a fixed value since the rehosted software will execute in a uniprocessor environment.

* Move word: Trap the operation code and perform the equivalent move operation with move characters (MVC). Alternatively, all instances of MVW could be replaced in the source code with equivalent MVC instructions.

4.2.3 NAS Monitor

The NAS monitor uses all of the special instructions except the display instructions and set PCI (which is not used in the NAS software). An instruction-compatible computer can support these special instructions in many ways. It will be important to develop support for these special instructions that is consistent with their frequency of use (static and dynamic), their function in a uniprocessor environment, and their effect on monitor performance. In addition to the special instruction support already described for DLY, LI and MVW, the following is one approach based on the FDP demonstration [FAAT81] for supporting the special instruction needs of the NAS monitor:

- Set configuration: Trap the operation code and update a virtual memory monitor table as necessary.

- Set address translation: Trap the operation code and modify the virtual memory page tables as necessary.

- Insert address translation: Change the operation code to a unused value to avoid the operation code conflict with move characters long (MVCL). Trap the reassigned operation code and access the virtual memory page tables as necessary.

- Load data address: The instruction need not be supported since this instruction is unique to the IOCE's and the IOCE functions will be replaced in the rehost system.

- Start IOCEp: The instruction need not be supported since the baseline rehost configuration would not have IOCE's. The NAS monitor would require modifications to achieve equivalent I/O control in the baseline rehost configuration.

- Store PS base resister: Trap the operation code and access the virtual memory page tables as necessary.

- Load PS base register: Trap the operation code and update the virtual memory page tables as necessary.

Four of the special instructions -- DLY, LDA, LI and MVW -- function somewhat differently in the IOCE. These differences are not expected to be a problem since the IOCE code would be replaced in the baseline rehost configuration (see App. G).

4.2.4 Display Channel

In the rehost system, each mainframe will be capable of providing all of the processing needs of a CCC and a display channel. The display channel support in a rehost system would be based on the 9020E software since the

Raytheon 730 software cannot be rehosted on a 9020 instruction-compatible computer. Note that a rehost system would replace both types of display channels - 9020E and Raytheon 730.

Most of the 9020E display channel software would be reused in the rehost system. The display device and configuration-dependent component of the display software will be replaced in the rehost system since the display interface would be the refresh buffers attached to the rehost system channels. The remainder of the display software will be reused without modification. Although no static instruction usage data are available for the 9020E software, the usage of non-standard instructions is expected to be limited to those used in the NAS application software and the four display instructions. The display instructions can be supported in many ways with an instruction-compatible computer. One approach would be to change the operation code for RPSB to a unused value and avoid the conflict with the operation code for compare logical characters long (CLCL) and then trap the operation codes for CSS, CVWL, RPSB and LC so that the equivalent functions could be emulated. This emulation would require careful investigation since the display instructions were originally implemented for performance needs.

The monitor currently used in the display channel based on the 9020E is very similar to the NAS monitor used for the CCC. During the rehost process, the two monitors would be merged so that only one version of the monitor would be supported (and maintained) and monitor code could be shared between the virtual processes in the rehost system.

4.2.5 Summary

The effects of rehosting the NAS application software have been shown to be minimal. That is, radar input processing would be revised and the remaining NAS application subsystems would not be changed as long as the VM monitor is augmented to support the 9020 special instructions. The display channel software would require modifications to accomodate the change from display element memory to the display buffers. The remainder of the display channel software is expected to be reusable with the VM monitor providing the equivalent function that the 9020 special instructions provide. The NAS monitor would require some modification as part of the rehosting process. That is, those parts of the monitor that support local devices, error analysis, reconfiguration, and startup would require revision. The remainder of the NAS monitor would be reusable as long as the VM monitor supports the 9020 special instructions.

Engineering estimates for the impact of rehosting the NAS software have been prepared for cost estimating purposes and are listed in the cost chapter (Table 5-1). These derived software costs include the costs to configure and augment the VM monitor.

4.3 9020 Environment

4.3.1 Introduction

The 9020 hardware and NAS monitor provide an operating environment for the NAS application software. Part of this environment is provided by the instruction-compatible computer and the support for the special instructions. The remainder of this environment would be provided by a combination of modifications to the NAS monitor and the services provided by a virtual machine monitor. The environment problem areas are:

- Memory usage,
- Timer usage and synchronization,
- Program status word (PSW) format,
- Devices and channel program usage,
- Diagnose and error analysis.

4.3.2 Memory Usage

There are several issues related to memory usage in the 9020 system--page zero, storage keys, immediate instructions and memory size. Page zero

(4096 bytes starting at byte 0) has many permanent storage assignments for system functions such as initial program loading, interrupt processing, I/O initiating, interval timer processing, and diagnostic logging. In a multiprocessor computer system, it is essential that this page for each active processor be relocated to a unique memory location. While the relocation is not essential for a uniprocessor environment, the NAS monitor assumes that the page zero will be relocated and refers to that relocated page in an absolute manner. One approach for resolving this problem is for the virtual memory monitor to use the parameters of SPSB and LPSB instructions to modify the page tables for the virtual memory associated with the NAS monitor.

Storage keys represent one mechanism for protecting memory in multiprogramming environments and they have been used to support the NAS software. Since some storage (Compool tables) is shared by several subsystems, the protection mechanism must allow access to shared storage. In the 9020 system, one storage domain was made accessible by all other storage domains. This sharing feature is not supported in any instruction-compatible computer. The problem could be resolved by micromanaging the storage keys in the virtual machine monitor. In the event that operational experience with the usage of storage keys represented a performance problem, then the rehost computer hardware could be modified (with an additional cost) to support the NAS usage of storage keys.

Three of the immediate instructions (and immediate, or immediate, and exclusive or immediate) operate with a fetch, modify, and store sequence that could cause undefined results in a multiprocessor configuration if one processor were to store a value into the same location that another processor had fetched a value from but before it had stored the value. In a multiprocessor environment, it is essential that these instructions execute in an atomic manner. This atomic execution would not be necessary for NAS application software in a uniprocessor environment since there would not be a competing processor. Interrupt processing in either multi or uniprocessing environments could interfere with the execution of immediate instructions. However, this interference would represent a software logic defect and not a rehosting problem.

The baseline rehost configuration will provide sufficient physical memory so that the NAS application software, the NAS monitor, and the virtual machine monitor can remain memory-resident at all times. Hence, disk buffering of PE's and flight plan tables as well as the CE overhead for the disk buffering will be eliminated and should result in a reduction of the average response time for ATC services that were supported by buffered PE's amd tables. The potential improvement in average response time may be limited by the internal queues for non-reentrant PE's (see 3.3.3).

4.3.3 Timer Usage and Synchronization

There are several timing considerations that must be analyzed as part of the change in NAS software environment. The NAS software needs both an interval timer and a time-of-day (TOD) timer which are currently provided by a 60 hertz decrementer at location 80 in page zero and a coded time source (CTS), respectively. Both of these timer needs should be supported by modifying the NAS monitor to directly access the timer support provided by the virtual machine monitor. Simulating the current timers with indirect access to the virtual machine monitor timer would result in accuracy problems.

The differences in processor and memory cycle times for the baseline rehost configuration as compared to the current 9020 processor and memory cycle times may result in some synchronization problems which have not been detected in the operation of the NAS software in the 9020A and 9020D configurations. That these differences would turn out to be a problem, however, seems remote; the rehosting contractor, nevertheless, should be aware of the possibility that a problem exists.

Another synchronization issue involves the use of write direct for interprocessor communication in a multiprocessor configuration. In the uniprocessor environment of the baseline rehost configuration, the write direct may be resolved by doing nothing since the current NAS monitor is capable of operating with one processor. Alternatively, the usage of write direct in the NAS monitor could be reviewed and the software revised as necessary. The companion instruction, read direct, is not an issue since it is not used in the NAS software.

4.3.4 PSW Format

The PSW format for the 9020 system is very similar to that for the standard IBM 360 computers with some unused bits in the interrupt code field assigned so that the additional 9020 channels could be addressed. Since the channel address problem has been resolved differently for IBM 370 computers and their equivalents, a PSW format difference would exist between the 9020 system and any replacement computer. This difference can be resolved in two ways:

- Allow the virtual memory monitor to translate the PSW format.

- Modify the NAS monitor so that all references to the PSW use the standard format. The changes would effect all uses of load PSW (LPSW) and set system mask (SSM) as well as part of the support for the supervisor calls (SVC).

4.3.5 Devices and Channel Program Usage

Since all of the local devices would be replaced as part of the baseline rehost configuration, all of the device support routines and procedures for accessing these devices will either be revised or replaced. The virtual machine monitor should provide all the necessary device support routines with minimal need for modifications. The channel programs in the I/O management subsystem and the I/O device-dependent code subsystem would require revision to accommodate the new devices. In the event that channel programs have been embedded in other parts of the NAS software, then these channel programs would have to be located and revised (or the adjacent code modified to use the standard I/O subsystems). In addition, those subsystems that are dependent upon device characteristics (for example, disk tracks and cylinders or tape density) would require modification.

The radar input is currently supported within an IOCE using an open loop channel program. In the rehost system, the radar input would be preprocessed by the radar input line multiplexor (see App. G) and presented in a blocked record format to the mainframe.

The interface between the mainframe in the baseline rehost configuration and the display generators would be a display buffer (see App. G). The refresh buffer would be capable of supporting the high data transfer rates required by the display generator and allow periodic access by the mainframe to update the display data.

4.3.6 Diagnose and Error Analysis

The diagnose instruction provides assistance in sorting out hardware problems. The functional operation and data values returned for this instruction differs for nearly every instruction-compatible computer and even for the 9020A and 9020D computers. Diagnose is an important part of the element error analysis and configuration and the I/O error analysis subsystems. Hence, these monitor subsystems would require careful analysis and revision not only to incorporate the rehost version of diagnose but to accommodate the baseline rehost configuration which is significantly different from the current 9020 configuration.

Another aspect of the error analysis is the requirement that the operation of the system be resumed as soon as possible after a failure has been detected and resolved. An essential part of resumed operations is to provide a valid copy of critical data values without resorting to complete reconstruction in the event of a detected compromise to the active database. In the present system, critical data values are written to disk, on a periodic basis (30 second interval) to facilitate database restoration.

In the rehost system, the reconfiguration process in the event of a failure would result in the transfer of active status to the "stand-by" processor. This transfer of status would be completed within 2 to 5 seconds since all of software in the "stand-by" processor would always be initialized awaiting access to the most recent set of critical data values. Note that a complete initiation of the rehost system starting with the initial program load for the virtual machine monitor would require at least 5 minutes.

In order to better support the reconfiguration process in the rehost system, the critical data values should be saved more frequently, perhaps with a 3 to 5 second interval. In addition, the range of critical data values saved should be reviewed in order to identify additional data values that would allow faster resumption of automated processes.

4.3.7 Support Software

The software support tools for the NAS system should be reusable in the rehost environment. Only the performance monitoring tools would require changes to reflect differences in timer support and device configurations. In particular, the high resolution timer (HRT) tool would no longer require a dedicated processor in order to generate high resolution timer intervals.

4.3.8 Summary

This section has considered a range of issues relating to the 9020 environment. The conclusion is that these issues can be resolved within the context of the VM monitor and the interaction between the NAS monitor and the VM monitor.

REFERENCES

[FAAT81] FAA Technical Center, Informal discussion with ACT-700 staff about their experiences in rehosting the Flight Data Processing component of the NAS application software on an IBM 4341 computer system operated with the VM/370 operating system. Atlantic City, NJ, May 1981.

[IBM73] IBM 9020D and 9020E System Principles of Operations, January 1973.

[IBM75] IBM SPAR 66, A3D2.2 RBB Instruction Sequence Scan Report, August, 1975.

5. COST

5.1 Introduction

The purpose of this chapter is to estimate the cost of rehosting the NAS software on instruction-compatible machines. This cost covers the development, acquisition, and operation of the new system. In estimating this cost six principles are followed.

First, the goal is to estimate the cost of instruction-compatible replacement relative to the cost incurred under the status quo. That is, the baseline against which cost is measured is the hypothetical situation in which the current system continued to operate through the rest of this decade. In other words, what is estimated is the change in the cost of providing en route air traffic control services that would result if instruction-compatible replacement were adopted.

Second, because the rehosting problem is not completely understood, the estimates in this chapter should be thought of as first approximations rather than as definitive. The goal is to give plausible estimates of what the cost of rehosting might be, but further study would be needed before one could have a high level of confidence in the cost estimates. The FAA personnel who provided the basic information used in this chapter operated under the understanding that what was needed was a reasonable first approximation and that they would be contacted again if a more accurate approximation was needed.

Third, a conservative approach is used in estimating the costs to make sure that the cost of rehosting is not underestimated; whenever there is doubt about a particular cost, a higher figure is chosen. Therefore, the cost estimated in this chapter can be thought of as an upper bound; effort has been made to make this upper bound as tight as possible.

Fourth, the procedure followed in this section is to spell out the basic data and the assumptions that are used to produce the cost estimates. It is not claimed that the data and assumptions are precise and perfect; all that is claimed is that the data and assumptions used reflect our understanding of the problem at the time this report was written. Every effort has been made to make clear what data and assumptions are used; the reader who has better assumptions or data should have no trouble with re-doing the calculation and producing his own estimates.

Fifth, it is assumed that the computers are replaced at all twenty centers. Appendix F discusses the case where replacement only occurs at some of the centers.

Sixth, all cost estimates are in 1981 dollars. No attempt has been made to estimate how these costs will change over time.

If instruction-compatible replacement were undertaken, the change in the cost of providing en route air traffic control services would fall into five broad categories:

- software: this includes the development and testing of new software and its integration with the old software and the new hardware;

- hardware: this includes the development, testing, and acquisition of the new hardware;

- maintenance cost: this includes the expenditure on personnel and parts made in order to maintain and support the system once it is in operation;

- transition cost: this includes the cost of remodeling needed to prepare the site for installation, of special hardware needed only for the transition period, and of training and other personnel costs.

- Program management and support cost: this includes the cost incurred by the FAA in administering the procurement.

Each of the costs will now be discussed.

5.2 Software Cost

The main advantage of replacing the 9020's with instruction-compatible machines is that the current NAS software would then be used on the new machines, and a wholesale rewriting of the software could be avoided. Nevertheless, some changes in the current software would be needed for the reasons discussed in Chapter 4, including the need to analyze and reconfigure new hardware when there is a failure, to handle new peripherals, and to deal with instruction set differences. An estimate of the money and time needed to carry out these changes in the NAS software will now be given.

Table 5-1 shows the NAS software subsystems and the size (in words) of each. HH Aerospace, after studying the rehosting problem, has estimated both the percentage of the words of code in each module that would be affected by rehosting and also the difficulty involved in dealing with this code; these estimates are shown in Table 5-1. It should be stressed that while the application code will be affected, it is not expected that it will be changed; the problems described in Ch. 4 will be taken care of by some method other than changing the application code. While this code will not be changed, it will have to go through testing and integration. Not shown in Table 5-1 are the changes that must be made to the virtual machine monitor to adapt it to the baseline rehost configuration and to modify it to support the handling of the non-standard instructions. It is estimated that 25,000 of the 500,000 words in the virtual machine monitor would need to be redesigned and recoded.

Software development and testing cost has been estimated with the PRICE S software cost estimation model. The estimation has been carried out by The Analytic Sciences Corporation (TASC) and is documented in a forthcoming report [TASC81]; the reader is referred to that report for details. Table 5-2 shows some of the assumptions used and Table 5-3 shows how the code is classified.

TABLE 5-1: ESTIMATED PERCENTAGE OF THE NAS SOFTWARE AFFECTED BY REHOSTING

Subsystem	Size	Affected	Difficulty
A. CCC Application Code			
Preliminary processing	18,692	20%	average
Flight data processing	53,212	10%	average
Route conversion	36,462	10%	average
Disk storage applications	22,732	20%	difficult
Posting determination	33,260	10%	average
Flight status alerts	46,872	10%	average
Inquiry processing	62,296	10%	average
Supervisory and interfacility	36,648	10%	difficult
Hardware error processing	6,934	20%	difficult
Track data processing	41,259	10%	average
Display channel outputs	48,652	20%	difficult
Real-time quality control	6,032	10%	average
Radar processing and tracking	31,400	50%	difficult
Flight plan analysis	2,562	10%	average
B. Monitor Code			
Startup/startover management	5,260	50%	difficult
Element error analysis and config.	10,980	100%	very difficult
Input/output management	4,046	20%	difficult
Input/output error analysis	2,566	100%	difficult
Program element control	1,256	20%	average
Program element synchronization	1,266	30%	difficult
Storage and communication mgmt.	2,432	20%	difficult
Man-machine communication	17,096	20%	difficult
On-line data recording services	6,900	50%	difficult
On-line test tools	18,688	50%	difficult
Contents supervisor	10,145	20%	difficult
Input/output device dependent code	11,542	100%	very difficult
C. Miscellaneous Code			
DCC	78,000	50%	difficult
CDC	48,700	0%	
DARC	53,000	0%	
MDM	2,000,000	0%	
NOSS	323,000	5%	average
Op support	1,000,000	0%	

Source: Size - [PDS80, Sec. 5.1.1]

TABLE 5-2: SOFTWARE DEVELOPMENT COST ASSUMPTIONS

CHARACTERISTIC	ASSUMPTION(S)
Type of System	MIL-SPEC ground-based aircraft control system.
Hardware Effects on Software	Capacity problems not anticipated. Response time problems not anticipated. Software can support all hardware interfaces in system.
Labor Costs	Per man-month labor costs are taken to be: \$6968 (design) \$5798 (implementation) \$5829 (test & integration)
Types of Software	The software to be developed consists of: application software (average difficulty) application software (difficult to develop) monitor software (average difficulty) monitor software (difficult to develop) monitor software (very difficult to develop) NOSS software (average difficulty)
Secondary Costs	15% of labor costs
Escalation	All costs in constant 1981 dollars
Integration of Software Into System Level Configuration	Typical level of integration effort anticipated.

Source: [TASC81]

TABLE 5-3: CLASSIFICATION OF THE CODE TO BE MODIFIED

NAS SOFTWARE TYPE	LEVEL OF DIFFICULTY	PRICE S APPLICATION CLASS
Application	Average	Real-Time Command & Control
Application	Difficult	Interactive Operations
Monitor	Average	Operating Systems
Monitor	Difficult	Operating Systems
Monitor	Very Difficult	Special class with an application class value 10% greater than the Operating Systems value.
NOSS	Average	Data Storage and Retrieval

Source: [TASC81]

TABLE 5-4: ESTIMATES OF THE SOFTWARE DEVELOPMENT AND TESTING COST

	Estimated Cost (millions)		
Category of Software	Low	Best	High
On-line	$2.520	$3.330	$3.650
VM modification	1.540	2.070	2.250
Support	0.330	0.420	0.450
Total	$4.390	$5.820	$6.350

Source: [TASC81]

AD-A118 828 H H AEROSPACE DESIGN CO INC BEDFORD MA F/G 9/2
A TECHNICAL ANALYSIS OF REHOSTING THE NATIONAL AIRSPACE SYSTEM --ETC(U)
OCT 81 K BRITTING, A COCANOWER, A GUPTA DT-FA01-80-Y-10595
UNCLASSIFIED HHA-81-4 FAA-AP-82-2 NL

2 of 2

END
DATE
FILMED
9 82
DTIC

Table 5-4 shows that the estimate of the software development and testing costs ranges from $4.390 million to $6.350 million, depending on the assumptions made. The best estimate is $5.820 million. This table breaks the cost down into the cost of changing the on-line software, the support software, and the virtual machine monitor. This cost estimate covers the development, testing, and integration of the new software; once this process is completed, the system is ready to be installed and tested at the FAA Technical Center.

5.3 Hardware Cost

The hardware cost that would be incurred under rehosting falls into the categories of mainframe cost, peripherals cost, special hardware cost, and system testing cost. Each of these will now be discussed. These costs are drawn mainly from a forthcoming TASC report [TASC81].

Mainframe cost. The two leading mainframes that are candidates for rehosting are the Amdahl 470/V7 and the IBM 3033U. Table 5-5 shows the cost of the mainframe and associated hardware that would be borne if there were rehosting. This table and the next are based on list prices. Each mainframe is assumed to have 8 megabytes of memory and 12 channels. The cost for each center is estimated to be about $4.3 million if the 470/V7 is selected and about $6.6 million if the 3033U is selected. (If the 470/V8 rather than the 470/V7 were selected, the added cost at each site would be about $300,000.)

One advantageous aspect of rehosting should be pointed out. Since off-the-shelf mainframes are used, this means that the processor can be upgraded if this proves desirable. For example, a V7 can be field-upgraded to a V8 at a cost of $250,000 [AMDA81]; this yields an increase in processor capacity which is estimated by one source to be 7 percent [LIAS80, p. 104] and by another to be 23 percent [HENK81, p. 14]. The ease of upgrading means that the FAA can avoid being pushed into overbuying by the uncertainty over how much processor capacity is needed.

TABLE 5-5: MAINFRAME ACQUISITION COST AT EACH CENTER

Amdahl	Unit Price	Number	Total
470/V7*	$2,125,000	2	$4,250,000
Channel to Channel Adapter	32,500	2	65,000
Two byte Interface	1,400	2	2,800
Total			$4,317,800
IBM			
3033U A09	$2,376,000	2	$4,752,000
Extended Addressing	93,900	2	187,800
Extended Control Store	24,800	2	49,600
Data Streaming	40,000	2	80,000
3033 Extension	35,000	2	70,000
MPU	287,000	2	574,000
Power/Coolant Unit	228,000	2	456,000
Console	192,000	2	384,000
Total			$6,553,400

*price includes power, cabinet, and console

N.B. Each mainframe is assumed to have 8 megabytes of main memory and 12 channels.

Source: [TASC81]

TABLE 5-6: PERIPHERALS ACQUISITION COST AT EACH CENTER

Item/Model	Unit Price($)	Number	Total
Mag. Tape 3240	24,190	4	$96,760
Mag. Tape 3803 Controller	36,815	2	73,630
Disk 3350	40,000	4	160,000
Disk 3380 Controller	97,650	2	195,300
Line Printer	51,130	2	102,260
L.P. Controller	17,685	2	35,370
I/O Switch	79,620	2	159,240
Total			$822,560

Source: [TASC81]

TABLE 5-7: ENGINEERING COST FOR THE SPECIAL HARDWARE

Special Hardware	Engineering Cost
RIN Line Multiplexor	$150,000
Refresh Buffer	150,000
Total	$300,000

Source: [TASC81]

TABLE 5-8: SPECIAL HARDWARE ACQUISITION COST PER CENTER

Special Hardware	Unit Cost	Units per Center	Total per Center
RIN Line Multiplexor	$ 3,500	25	$ 87,500
Refresh Buffer	10,000	15	150,000
Cabinet, Power and Connectors	1,000	2	2,000
Total			$239,500

Source: [TASC81]

TABLE 5-9: SUMMARY OF THE HARDWARE COST

Engineering cost	$ 300,000
Acquisition cost per center*	
Amdahl 470/V7	5,379,860
IBM 3033U	or 7,615,460
System testing	6,300,000

* includes acquisition of mainframe, peripherals, and special hardware

Source: Tables 5-5 through 5-8

Peripherals cost. Table 5-6 shows that $0.8 million is the cost at each center of replacing the magnetic tape units, the disk units, the line printer, and the control units.

Special hardware cost. The special hardware required by rehosting is the RIN line multiplexor and the refresh buffer as described in App. G. For each piece of special hardware there is a one-time engineering cost for the system that covers design, development, testing, and software. This one-time engineering cost is estimated to be $300,000, as Table 5-7 shows. The special hardware acquisition cost for each center is estimated to be $239,500. Table 5-8 shows the number of units needed at each center and the unit cost.

System testing cost. After the hardware and software are developed and tested by the contractor, the FAA will test the complete system at the FAA Technical Center and then at an en route center. For convenience, this cost is included here. The testing process is expected to take 15 months (see Ch. 7) and to cost $5.0 million per year, a figure provided by the FAA. Therefore, the total system testing cost is estimated to be $6.3 million.

Summary. The hardware cost is summarized in Table 5-9. The engineering cost for the special hardware, which is incurred only once for the system as a whole, is estimated to be $300,000. The cost of acquiring the mainframe, peripherals, and special hardware for each center is estimated to be either $5.380 million (if 470/V7's are procured) or $7.615 million (if 3033U's are procured). The system testing cost, which is incurred once, is estimated to be $6.300 million.

5.4 Maintenance Cost

5.4.1 Introduction

Maintenance and support of the 9020's is carried out by personnel at the ARTCC's, the FAA Technical Center, and the FAA Depot. If there is a failure, the problem is diagnosed by personnel at the ARTCC, perhaps assisted by Technical Center personnel. If the failure is a hardware failure, the faulty part, once located, is replaced. It is repaired at the center or, if the repair is complex and the part costs more than $300, it is sent to the Depot under the exchange and repair program. The Depot then sends a good part from its stock to the ARTCC and repairs the faulty part if possible and adds it to its stock. The Depot is responsible for providing virtually all spare parts to the centers.

This maintenance strategy might be changed in two ways under rehosting. First, the FAA might find it advantageous to diagnose problems by using a telephone link with a remote diagnostic center. Second, because of the use of large scale integration, it is not really feasible for the FAA to repair failed cards; repair of these cards would require very elaborate facilities and would probably be done by the manufacturer. Nevertheless, for purposes of estimation, it is assumed that the rehost system is maintained in the same way as the current system.

The amount by which the maintenance and support cost would change under rehosting will now be estimated. This cost is divided into the two categories of personnel and parts.

5.4.2 Personnel

The expected change in the wages paid to maintenance and support personnel will now be estimated; costs of training these personnel are considered under transition cost in Sec. 5.5. Since the changes made to the software would not greatly alter its size or structure, it is assumed that there would be no change in the software maintenance cost.

It is expected that there will be a decrease in the hardware maintenance cost because of the technological progress that has occurred since the 9020's were purchased. Not only is modern hardware more reliable, as Ch. 2 points out, but when there is a failure it is easier to diagnose and repair. Also, since the replacement system will have fewer components than the old, there will be a lesser need for specialists.

TABLE 5-10: ANNUAL REDUCTION IN THE COST OF HARDWARE MAINTENANCE PERSONNEL AT A TYPICAL ARTCC

Year	Reduction
First	$ 0
Second	137,431
Third	274,861
Fourth and later	412,292

The magnitude of the reduction in the personnel cost of hardware maintenance is estimated in the following way. A recent report commissioned by the FAA estimates that at a typical ARTCC there are the equivalent of 33.9 full-time Airway Facilities Service personnel working on automation [ASI80, p. 4-4]. Assume that the average grade is a GS-13, step 4, with an annual salary of $35,252; increase this by 15 percent to $40,540 to cover benefits and overtime. This gives a government outlay of $1,374,306 at a typical ARTCC.

It is assumed that during the first year that an ARTCC has the new system, there will be no reduction in the cost of hardware maintenance personnel because of the frictions of transition. It is assumed that there is a 10 percent reduction in each of the second, third, and fourth years; therefore, the long-term reduction is 30 percent. (The Airway Facilities Service has stated that 30 percent is a reasonable figure.) The dollar amounts that would be saved per center each year are shown in Table 5-10.

The long-term reduction of 30 percent is chosen since it is a conservative figure that appears not to overestimate the savings that rehosting would provide. Though the number of hardware failures is expected to fall by 60 to 90 percent, this lower figure of 30 percent is chosen for two reasons. First, the reduction in personnel is less than the reduction in failures because of the need for specialists. For example, even if there is a greatly reduced number of memory failures, it is still necessary to have a specialist who can deal with memory. Second, the figure of 33.9 full-time personnel is slightly too large since it includes AF personnel who maintain the software in the display channel, which is not relevant.

Since the number of relevant hardware maintenance personnel at the Technical Center and the Depot is small, no reduction in cost at these organizations is estimated. There are about 25 AF personnel at the Technical Center. The chief of the Engineering and Production Branch at the Depot estimates that the equivalent of between 2 and 3 full-time technicians work on 9020 parts in the exchange and repair program.

5.4.3 Parts

The change in the cost of replacement parts that would result from rehosting can be divided into the start-up costs and the annual cost.

Start-up cost. Information on the start-up costs that would be incurred in laying in an initial inventory of replacement parts and in meeting other front-end requirements was provided by the chief of the NAS Project and Provisioning Section of the Depot. These costs are expressed as a percentage of the hardware acquisition cost for a single center. These percentages are based on rules of thumb and on experience with other systems, not on a study of the rehosting problem. Therefore, these percentages should be thought of only as first approximations. There are start-up costs both for the Depot and for each center.

For the Depot there are two start-up costs. First, 6 percent of the hardware cost at one site is assumed to cover (a) documentation on engineering and provisioning, including engineering drawings and all other engineering specifications (as required by FAA-G-1210d [FAA78]); (b) training on how to use the testbed; (c) development of equipment to troubleshoot the system. Second, 10 percent is needed to purchase a testbed, which is the hardware needed to test parts that have been repaired to insure that the repairs have been made properly. The start-up cost at the Depot, then, is

$$\$5.380 \text{ million} \times 0.16 = \$0.861 \text{ million.}$$

if V7's are procured or

$$\$7.615 \text{ million} \times 0.16 = \$1.218 \text{ million}$$

if 3033U's are procured.

For each center there is a start-up cost of 24 percent. This goes for spare parts, some of which are stocked at the center and some at the Depot. This figure breaks down into 4 percent for parts common, i.e., parts that can be ordered from a vendor's catalog, and 20 percent for parts peculiar, i.e., parts that are not parts common. Therefore, the start-up cost at each center is

$$\$5.380 \text{ million} \times 0.24 = \$1.291 \text{ million}$$

if V7's are procured or

$$\$7.615 \text{ million} \times 0.24 = \$1.828 \text{ million}$$

if 3033U's are procured. Therefore, the initial stock of spare parts for all 20 centers is $25.820 if V7's are procured and $36.560 if 3033U's are procured. It should be stressed that this figure of 24 percent is very conservative. Moreover, it does not take into account the dramatically improved reliability of the new system (as discussed in Ch. 2). Therefore, it is probable that the estimate of the initial spare parts cost is much too high.

<u>Annual cost.</u> According to the chief of the General Materiel Section of the Depot, the FAA has a contract with IBM under which the FAA buys the replacement parts needed for the 9020A,D, and E systems. For the last few years the annual cost of the spare parts for these systems has hovered around $950,000. Similar information on the cost of replacement parts for the Raytheon 730 is not available, so the following rough approximation will be used. There are 25 9020 A,D, and E systems and 15 Raytheon 730 systems. Therefore, assume that the cost of parts for the 730's is 15/25 that of the cost of parts for the 9020's, i.e., $570,000. This means that the total annual cost of replacement parts for the CCC's and display channels is $1,520,000. (This cost is estimated to be $2.3 million in a report on maintenance cost prepared for the FAA [ASI80, p. 6-6]. This estimate, however, seems to be based on less reliable information, and it is not used here.)

It is assumed that with a new system the annual expenditure on parts will fall by two-thirds, i.e., by $1,013,333. Even though the actual parts usage will, it is thought, fall by somewhat more than two-thirds , this lower figure is used to make sure that the saving is not overstated and to allow for the possibility that the current parts usage of the Raytheon 730 has been overstated.

Table 5-11 summarizes the effect that rehosting would have on the expenditure on cost. Initially there would be a one-time cost of either $26.681 million if V7's are procured or $37.778 if 3033U's are procured; this would lay in a stock of replacement parts and set the Depot up so it could deal with the new system. There would, however, be a saving of $1.013 million each year for the system because the new system would require fewer replacement parts.

5.4.4 Summary of the Estimated Annual Savings in Maintenance Cost

The time profile of the savings on personnel and parts will now be considered. It is assumed that the new system will go into operation at the 20 ARTCC's over a period of 2 years (see the procurement schedule in Ch. 7). Therefore, it is here assumed that 10 systems go into operation in the first year and 10 in the second. The yearly savings are shown in Table 5-12 and will now be explained.

Personnel cost. For any one center it has been assumed that the cost of hardware maintenance personnel will not change in the first year of the new system but will then decline by 10 percent each of the next three years for an eventual annual saving of 30 percent of the estimated current figure of $1.374 million per year. Consider the 10 systems installed in the first year. There is no saving in the first year. The saving in the second year is $1.374 million, i.e., $1.374 million x 10% x 10 centers. The savings in the third and fourth years are then $2.748 million and $4.122 million. For the 10 systems installed in the second year, the savings are also $1.374, 2.748, and 4.122 million, realized in the third, fourth, and fifth years respectively. These two streams are then added together to obtain the total personnel saving per year, which is shown in Table 5-12.

Parts. Since the annual saving in parts with the new system is estimated to be about $1.013 million per year, half this amount is saved the first year when half the systems are in operation, and the full $1.013 million is saved in subsequent years. These figures are shown in Table 5-12.

5.5 Transition Cost

Transition cost covers all the costs that are incurred because of the switch-over to a new system and can be divided into four categories: remodeling cost, special hardware cost, extra personnel cost, and training cost. Each of these categories will now be discussed.

Remodeling cost. If the broadband is removed from the centers in 1984 as planned, then there should be sufficient floorspace to comfortably house the old and new systems simultaneously; this means that no major construction would be needed [MULL81, pp. 39-40]. The cost of remodeling is estimated by the FAA to be one million dollars per center.

Special hardware cost. There will be a need for hardware that will be thrown away after the transition period, e.g., extra cables and switches. It is expected that this would be minor, so no cost is assigned.

Extra personnel cost. This cost refers to the extra personnel that might be needed to help make the transition to the new system. These extra personnel might be needed just before replacement when the heavy training schedule has temporarily depopulated a center. The extra personnel might also be needed during the period of parallel operation; with two different systems operating, the center's normal staff might be overtaxed. (This topic is further discussed in Chapter 6). It would be premature to specify how the transition would be made and how many extra personnel would be needed. To serve as a round figure representing the cost of extra personnel, $200,000 per center is chosen.

TABLE 5-11: CHANGE IN THE EXPENDITURE ON REPLACEMENT PARTS DUE TO REHOSTING

One-Time Costs		Cost (millions)
V7:	At the Depot	$ 0.861
	At 20 Centers	25.820
	Total	$26.681
3033U:	At the Depot	$ 1.218
	At 20 Centers	36.560
	Total	$37.778
Annual Cost for the System		($1.013)

N.B. A figure in parentheses is a reduction in expenditure.

TABLE 5-12: ANNUAL MAINTENANCE COST SAVING PROVIDED BY REHOSTING (millions)

	Year				
	First	Second	Third	Fourth	Fifth and after
Personnel	$0.0	$1.374	$4.122	$6.870	$8.244
Parts	0.506	1.013	1.013	1.013	1.013
Total	$0.506	$2.387	$5.135	$7.883	$9.257

Training cost. If confronted by a new system, those who operate and maintain the system would require training. The cost of this extra training that would be necessitated by rehosting will now be estimated. This training cost can be divided into the cost of developing the new courses and the cost of teaching them.

The cost of developing the new courses is estimated in the following way. Table 5-13 shows the relevant AF courses currently given at the Academy. These are the courses that would have to be replaced if there were rehosting. Omitted from this list are courses that would still be given under rehosting without significant change (e.g., courses on Jovial programming, on the applications programs, and on hardware that would be retained), courses that would not be replaced since the subject matter would not be relevant under rehosting (e.g., courses on the display channel hardware), and courses not needed because there will be only one system for the CCC and display channel rather than three as at present.

The courses that would be given on the new system would probably be structured differently from the current courses; nevertheless, the courses listed in Table 5-13 will be used as a rough guide to what the new courses might look like. (AT courses are not considered since there are only a few of them and since they would be relatively untouched by rehosting.) These courses together last a total of 81.8 weeks, or 3,272 hours. This figure will be increased to 3,500 hours to allow for any new courses not captured in Table 5-13.

The cost of developing these new courses can now be estimated. The chief of the Automation Section of the Airway Facilities Branch of the Academy has provided the rule of thumb that the ratio of development hours to course hours is 30 to 1. Therefore, the number of hours needed to develop the new courses is

$$3,500 \times 30 = 105,000.$$

The Budget Division of the Aeronautical Center puts a cost of $16.94 on each hour of productive time spent by AF instructors. (This is calculated by

TABLE 5-13: RELEVANT AIRWAY FACILITIES COURSES OFFERED AT THE ACADEMY

Course Number	Length (weeks)	Cost per student per week	Title
43458	8	$119	IBM 9020 System Familiarization and BAL Programming
43459	8	101	IBM 9020 Input-Output Equipment
43460	6	125	IBM 9020 A/D PAM and System Control
43462	20	143	IBM 9020D/E Processing
43437	4	123	IBM 2314-A1 Direct Access Storage Facility
43432	3	136	System Maintenance Monitor Console for Technicians
43468	5	112	OS-360 and DASF Programming Techniques
43469	10	165	FDP and Monitor for Systems Performance Specialists
43470	6	126	NAS Operational Program for Engineers
43471	5	108	NAS System Interface
43489/ 90/91	6.8	NA	CCC for Engineers
Total Weeks	81.8		

taking the average salary of AF instructors of $27,748, increasing it by 27 percent to $35,240 to allow for benefits, leave, and training, and dividing by 2080.) Therefore, the estimated cost of developing the new courses is

$$\$105,000 \times \$16.94 = \$1,778,700.$$

It would probably be the case that these courses would be developed largely by a contractor rather than the FAA.

The cost of <u>teaching</u> the new courses will be estimated by developing an equation that expresses the cost of training a student as a function of the length of his training and then by combining this with information on the amount of training that would be required.

The main source used to derive the cost of training a student is a document compiled by the Budget Division of the Aeronautical Center [AER081]. This document estimates the cost of teaching each course <u>per student</u>. The cost is divided into three components. The first cost is personnel compensation and benefits, i.e., the instructor's time; this includes both the time spent in the classroom and the time spent in preparation. This cost is based on the actual hours reported by the instructors. The second cost is supplies and course material such as manuals. The third cost is overhead to cover administration, buildings, and so forth; this cost would not be affected by rehosting and is, therefore, not included in the estimates of training cost. For example, for the 9020 D/E processing course, number 43462, the per student cost for personnel compensation and benefits is $2,859 and for supplies is $5, for a total of $2,864. Since this course is 20 weeks long, the cost per student per week is $143. The cost of other courses per student per week is similarly calculated and shown in Table 5-13. A weighted average of these costs is taken, where the weights are proportional to the length of the course, and the resulting average cost per student per week is $130.

The per diem rate for a stay of two weeks or longer is $31, so the per diem cost is $217 per week. This means that the average cost per week for the course and the per diem for a student is $130 + 217 = $347, which is rounded to \$350. The average round trip travel cost to Oklahoma City is assumed to be \$400. Therefore, we have the training cost equation

$$TC(w) = \$400 + \$350w,$$

which expresses the training cost for a single student as a function of the number of weeks w he is at school. This equation does not include the salary of the student since this cost would be incurred even if he were not attending the course.

The next task is to estimate the number of people who would require training and the amount of training that is necessary. To do this the relevant organizations will be examined.

Consider the en route centers. The two relevant organizations are the Air Traffic Service (AT) and the Airway Facilities Service (AF). As a first approximation the division of responsibility between these organizations is that AT maintains the software in the CCC and AF maintains all hardware as well as the software in the display channel. Typical though not invariable staffing patterns for AT and AF at a center are shown in Tables 5-14 and 5-15, respectively. Table 5-16 shows the AT staffing at the FAATC. (These staffing patterns were provided by FAA personnel. Table 5-15 is an almost perfect match with a similar table in [ASI80, p. 44], which, after subtracting out secretaries, shows 47 AF employees at each center.) At the FAATC AF has about 25 people; roughly 15 work on the 9020A, D, and E, and 10 work on the Raytheon 730. R&D, ARD-140, at the Technical Center has roughly 20 people; about 12 work on near-term enhancements and 8 on specifications for 9020 replacement. ACT-700, which operates and manages the 9020's at the Technical Center, has about 100 relevant employees. There are also a number of contractors at the center, but no costing of contractor training will be attempted.

Considering the duties of each organization and what would be affected by rehosting, estimates of the amount of needed training have been made and are shown in Table 5-17. This table shows for each organization at the en route centers and the Technical Center the estimated number of people who would need training and the average length of the training. The length of the training, when substituted into the training cost equation above yields the cost per student shown in Table 5-17. Multiplying this times the number of students yields the total training cost for each organization. It is seen that the estimated training cost is $370,400 at each center and is $630,000 at the Technical Center.

Table 5-17 also shows the cost of training the instructors at the Academy. The rule of thumb provided by the chief of the Automation Section at the Academy is that the ratio of hours needed to train the instructors to the hours of course time is 1.5 to 1. Therefore, the cost of training the instructors is estimated to be

$$\$3,500 \times 1.5 \times \$16.94 = \$88,935.$$

The final entry in Table 5-17 is the total cost of teaching the new courses during the transition. This cost of $8.127 million is obtained by adding the cost at each center times 20 to the costs at the Technical Center and at the Academy.

This completes the estimation of the cost of developing and teaching the courses that would be required by rehosting. Because of doubts about exactly what training would be required and about the accuracy of the rules of thumb, these estimates should be thought of as approximations rather than as precise estimates. Omitted from these estimates are the costs of training contractors and personnel in Washington and in the regional offices.

Summary. The various transition costs are summarized in Table 5-18, which shows that the estimated cost of remodeling the centers, paying the extra personnel needed during the transition, and developing and teaching the courses comes to about $36.906 million.

5.6 Program Management and Support Cost

If rehosting is adopted the FAA will incur administrative costs as it plans, reviews, oversees the procurement, and provides general support. The

TABLE 5-14: TYPICAL AIR TRAFFIC SERVICE STAFFING AT A CENTER

1	data systems officer
5	operations specialists (who monitor overall operation of the computer from the AT viewpoint)
4	adaptation specialists
4	testing specialists
7	programmers
21	Total

TABLE 5-15: TYPICAL AIRWAY FACILITIES SERVICE STAFFING AT A CENTER

5	system performance specialists
1	system performance officer
3	staff engineers or technicians in depth
7	computer operators
10	system engineers and assistant system engineers
20	technicians
46	Total

TABLE 5-16: AIR TRAFFIC SERVICE STAFFING AT THE FAA TECHNICAL CENTER

10	design team
22	production team
20	testing team
16	field support
2	documentation
6	supervisors
76	Total

TABLE 5-17: TRAINING COST NECESSITATED BY REHOSTING

	Number Requiring Training	Length of Training (weeks)	Cost per Student	Total Cost
At each Center				
AF	46	20	$7,400	$340,400
AT	12	6	2,500	30,000
Total				$370,400
At the Technical Center				
AF	20	32	$11,600	$232,000
AT	50	6	2,500	125,000
R & D	20	20	7,400	148,000
ACT-700	50	6	2,500	125,000
Total				$630,000
At the Academy				88,935
Total Cost of Teaching Courses				$8,126,935

FAA has estimated that this cost would over the life of the procurement amount to $41.2 million. Table 5-19 shows the breakdown by year. The tasks in each year can be seen by looking at the procurement schedule in Ch. 7.

5.7 Summary

The costs that are relevant to rehosting fall into two types. The first type is the front-end costs that are incurred initially to get the rehost system operational at all sites. Table 5-20 summarizes the initial cost estimates that have been made throughout this chapter; the total is $241.0 million if Amdahl 470/V7's are procured or $303.4 if IBM 3033U's are procured. (HQ, TC, and AC stand for FAA headquarters, the Technical Center, and the Aeronautical Center, respectively.)

The second type of cost is the annual cost of operating and maintaining the systems. The annual cost savings provided by rehosting are summarized in Table 5-21. Once the three year shakedown period is completed, the annual saving in personnel cost and parts cost is estimated to be $9.3 million.

The point about Table 5-20 to be emphasized is the dominance of the hardware cost. The hardware acquisition cost is about half of the total cost. Moreover, the next largest cost, the cost of the initial stock of spare parts, is closely tied to the hardware acquisition cost. The dominance of hardware acquisition means that efforts to hold down this cost can have a much bigger payoff than efforts to hold down other costs. Also, uncertainty over this cost dwarfs all the other uncertainties.

App. F discusses the idea of saving on the hardware acquisition cost by making a partial replacement, i.e., rehosting at some sites but keeping the 9020's at others. The conclusion is that if V7's are procured, the initial cost of rehosting falls from $241.0 million to $107.0 if there is replacement at 5 centers instead of 20; if there is replacement at 10 centers, the initial cost is $151.7. The long-term annual saving of $9.3 million falls to $1.6 million if there is replacement at 10 centers and is a $0.8 million annual increase if there is replacement at 5 centers. These

TABLE 5-18: SUMMARY OF THE TRANSITION COSTS
(millions)

Remodeling 23 sites	$ 23.000
Extra personnel at 20 sites	4.000
Course Development	1.779
Teaching	8.127
Total Transition Cost	$ 36.906

TABLE 5-19: PROGRAM MANAGEMENT AND SUPPORT COST (millions)

Year	Cost
1	$ 9.1
2	11.6
3	6.3
4	6.3
5	6.3
6	1.1
Total	$41.2

TABLE 5-20: INITIAL COSTS INCURRED BY REHOSTING (millions of dollars)

	Site					
	HQ	TC	AC	ARTCC's		Total
Software	5.8					5.8
Hardware						
Engineering	0.3					0.3
Acquisition						
V7		10.8	5.4	107.6		123.8
3033U		15.2	7.6	152.3	or	175.1
Testing		3.8		2.5		6.3
Maintenance						
Initial cost						
V7			0.9	25.8		26.7
3033U			1.2	36.6	or	37.8
Transition Cost						
Remodeling		2.0	1.0	20.0		23.0
Extra personnel				4.0		4.0
Developing courses			1.8			1.8
Teaching courses		0.6	0.1	7.4		8.1
Program management and Support	41.2					41.2
Total						
V7						241.0
3033U					or	303.4

TABLE 5-21: ANNUAL MAINTENANCE COST SAVING PROVIDED BY REHOSTING

Year	Saving (millions)
1	$0.506
2	2.387
3	5.135
4	7.883
5 and after	9.257

Source: Table 5-12.

estimates are incomplete since they do not reflect the inconvenience that would result from there being two different systems in the field.

Another suggestion for decreasing the cost is to follow an alternate, less finely tuned rehosting approach than that assumed in this report. This alternate approach would keep the present peripherals and make fewer changes in the software; this would entail a much higher VM overhead. This approach would cut perhaps $25 million off the cost (largely because the cost of new peripherals would be saved) and perhaps 12-15 months off the procurement schedule (since a much less extensive system development would be needed).

REFERENCES

[AERO81] "FY-1981 Schedule of Reimbursable Costs for FAA Academy Courses," prepared by the Aeronautical Center, FAA, n.d.

[AMDA81] "Authorized ADP Schedule Pricelist," prepared for the General Services Administration by the Amdahl Corporation, covering the period 3-4-81 through 9-30-81.

[ASI80] Automated Services, Inc., "Working Paper on NAS Automation Equipment: Operating Cost Estimates, FY 1978-1984," prepared for the Planning Requirements Branch, Office of Aviation Systems Plans, FAA, September 1980.

[FAA78] "Provisioning Technical Documentation," Federal Aviation Administration Specification, FAA-G-1210D, 1978.

[FAA80] "The FAA Plans and Programs for the Future Airport and Traffic Control System: Presentations at the Air Traffic Control Computer Replacement Industry Briefing," prepared by the FAA, December 2, 1980.

[HENK81] Tom Henkel, "IBM and the Plug-Compatibles," Computerworld, July 13, 1981, pp. 12-16.

[LIAS80] Edward J. Lias, "Tracking the Elusive KOPS," Datamation, November 1980, pp. 99-105.

[MULL81] D. F. Muller, "9020 Replacement Impact on ARTCC Environmental Systems," prepared by the staff assistant to the Associate Administrator for Engineering and Development, FAA, January 9, 1981.

[PDS80] "Program Design Specification: National Airspace System Air Traffic Control Computer Program, Volume II: Application Systems, Model A3d2.10," NASP-5105-16 (Vol. II), 14 April 1980.

[TASC81] The Analytic Sciences Corp., forthcoming report on the FAA's options for dealing with the en route computers, prepared for the Office of Systems Engineering Management.

6. TRANSITION

The FAA has established the requirement that when the existing en route computers are replaced, the transition from the old to the new system must be smooth and trouble-free so that safety is not jeopardized. A key ingredient in meeting this requirement is having a 90 day period of parallel operation of the two systems [FAA80, p.17]. This insures that there will be a proved, reliable back-up if there are any problems with the new system. If this somewhat intricate, non-standard transition is to be accomplished successfully, three types of problems must be dealt with:

- remodeling problems;
- technical problems;
- personnel problems.

Each problem area will now be briefly discussed.

<u>Remodeling problems</u>. Because of the need for parallel operation of the two systems, there must be enough room in the centers to accommodate both systems simultaneously. Once the direct access radar channel is field-tested and is operating normally, the plan is to remove the broadband radar from the centers; this will free up about 5000 square feet [MULL81, pp.39-40]. If this happens by 1984 as planned, then there will be sufficient space for the two systems and there will be no need for major construction. (This space, however, might be in an undesirable location such as a basement.) Since only remodeling would be required, no significant transition problems are expected. If, however, the broadband is not removed on schedule, then it is possible that there could be insufficient space, and there could be a need for major construction or for temporary shelters.

<u>Technical problems.</u> The technical problems posed by parallel operation fall into three areas. First, it is necessary to be able to feed the input signals into either (or both) of the two systems. These inputs include both external inputs (e.g., from radar) and internal inputs (e.g., from controllers). While all these signals have a relatively low data

transmission rate, the interface to these signals will be at the channel side of the PAM's and DAU's rather than at the termination of the input circuits since the PAM's and DAU's will be reused in the rehost system. The basic procedure for providing access to the PAM's and DAU's from the current and the rehost systems is to extend the current capability of the line controllers that allow them to be connected to more than one channel. After this multi-channel interface capability has been expanded, then the line controllers would be connected to the rehost system. The procedures and schedules for this shared access to the input signals must be carefully planned to ensure continuous operation of the entire system and minimize the disruption to each line controller as the shared access is implemented. The physical placement of new cables under the computer room floor will require careful planning as well since a large number of cables have already been located in the work area under the floor. While the display channel using the Raytheon 730 has a different line controller for the input circuits than the 9020D, this interface is not expected to be more complex than that for a DAU.

Second, if the rehost system fails, it is necessary that the 9020 system be able to take over and prevent a serious interruption in service. This can be done provided that the 9020 has access to a current data base. There are several schemes that could provide this current data base. In one leading scheme the inputs are fed into both the old and new systems, and both systems operate continuously. This means that if the new system fails, then the old system has a current data base since it has been maintaining it. The switch from the new to the old system would probably be done manually, so there would be a lag between the occurrence of a failure and the time that this failure is detected and the switch thrown. Once the switch is thrown, it would only be a matter of milliseconds until the old system comes on line and is providing full service. In short, it does seem possible to make a smooth transition from a failed system to the other system.

Third, it is necessary to feed output signals from either the current system or the rehost system to their appropriate destinations. These outputs fall into two groups—outputs for the display generators and outputs to low data transmission rate internal and external circuits. The transition procedure for the second group of outputs would be part of the input transition procedure since the same line controllers are used for input and output. Shared access to the display generators would be provided in a manner similar to that for the line controllers.

Personnel problems. The need to at all times have personnel at each center to maintain the system(s) leads to two possible problems. First, before replacement occurs, a large number of personnel at each center will require training, but this training must be scheduled so that enough personnel are left at the center to provide adequate support. Second, during the period of parallel operation there must be sufficient personnel to support both systems.

No plan to deal with these problems will be spelled out, but it is clear that these problems can be dealt with. For example, having the contractor hire and train extra personnel that would float from center to center as needed would be one possible way of dealing with these problems. The specific plan adopted should deal with a number of questions.

- When should the new system be installed at the Academy in Oklahoma City? If installed too late, it will not be available when it is needed for the initial surge of training.

- How will training be scheduled? If training is too early, skills will deteriorate before the new system is installed; if training occurs just before replacement, this might leave the center undermanned.

- How much of the training can be computer-based instruction that occurs at the centers?

These are all problems that, while not insuperable, do require careful planning.

Summary. It is seen that there are a number of complex problems that attend the transition period of parallel operation of the two systems. Though a heedless transition would run afoul of these problems, it seems fair to say that they can be successfully handled by advance thinking and careful preparation.

REFERENCES

[FAA80] "The FAA Plans and Programs for the Future Airport and Traffic Control System: Presentations at the Air Traffic Control Computer Replacement Industry Briefing," prepared by the FAA, December 7, 1980.

[MULL81] D. F. Muller, "9020 Replacement Impact on ARTCC Environmental Systems, prepared by the staff assistant to the Associate Administrator for Engineering and Development, FAA, January 9, 1981.

7. PROCUREMENT SCHEDULE

In deciding whether rehosting is a suitable method for extending the life of the current system, two questions arise. First, when will the current system need to be upgraded in order to avoid capacity problems? Second, could rehosting be accomplished fast enough to provide the needed upgrading? This chapter will not discuss the question of when and if upgrading will be needed; that question is being examined by other studies being conducted by the FAA. This chapter will, however, examine the question of how quickly rehosting could be accomplished. The goal is to estimate how much time elapses between the time the FAA issues the request for proposals (RFP) and the time that the rehost system is operating normally. This elapsed time is estimated by considering the six stages of the procurement that follow the issuance of the RFP; the duration of each stage is derived from discussions with FAA personnel.

First, potential contractors prepare proposals that spell out the approach to rehosting that the contractor plans to follow: 3 months.

Second, the FAA evaluates the proposals and awards a contract: 6 months.

Third, the contractor who is selected develops the software and hardware that his rehosting approach requires: 21 months.

Fourth, the contractor delivers and installs a system at the FAA Technical Center (FAATC), and this system is then subjected to full testing: 9 months.

Fifth, the system is installed at an ARTCC and fully tested and brought to the point where it is operational: 6 months.

Sixth, the system is installed and made operational at the remaining ARTCC's: 24 months.

Table 7-1 summarizes the stages and the elapsed times of this prudent, conservative procurement schedule. From the time that the RFP is issued until the first system is fully tested and operating normally at an ARTCC, 45 months (3 years, 9 months) elapses; from the decision to rehost until replacement is complete, 69 months (5 years, 9 months) elapses. This means that if the RFP were issued on 1 July 1982, then the first rehost system at an ARTCC would be operational on 1 April 1986; the rehost systems would be operational at all sites on 1 April 1988.

It is possible that the need for rehosting would be seen as urgent and that Congress would mandate that rehosting be accomplished as quickly as possible, with speed being achieved by cutting down administrative delays. This report will not speculate on how much the procurement schedule could be compressed under these circumstances.

It has been suggested that the mainframes could be leased rather than purchased to shorten the procurement cycle; this approach has three problems. First, since developing the system rather than acquiring the mainframes is the bottleneck in the procurement, leasing would not speed the procurement. Second, since three years is typically the break-even point for a lease and since these computers would probably be in place for more than three years, leasing would end up costing more than purchasing. Third, the user is typically not allowed to maintain the leased computers, and this would interfere with the FAA's providing the type of maintenance required by air traffic control.

One way to shorten the procurement schedule by perhaps 12-15 months would be to follow the alternate rehosting approach mentioned in the last paragraph of Ch. 5.

TABLE 7-1: THE PROCUREMENT SCHEDULE

Elapsed Time (months)	Stages of the Procurement
3	Industry: prepares proposals
6	FAA: evaluate proposals and awards contract
21	Contractor: develops software and hardware
9	FAA and contractor: test system at the FAATC
6	FAA and contractor: install and test system at the first field site
24	Contractor: installs systems at the remaining centers
69	Total

8. GROWTH POTENTIAL

8.1 Introduction

In order to minimize the future trauma of transitioning to a new system, the FAA has specified that any replacement system must be able to evolve smoothly over the next few decades. In particular, any replacement system must be capable of:

- Accommodating new hardware so that:
  - + The capacity of the system can be increased and the response times for all ATC related activities can be maintained at specified levels,
  - + New hardware technology can be integrated into the system in an evolutionary manner.
- Accommodating the evolution of ATC functions so that:
  - + Existing capabilities can be refined and extended,
  - + New capabilities can be added that automate more of the ATC process.

The ability of the rehosted system to meet the needs will now be discussed.

8.2 Hardware Growth Potential

The baseline rehost configuration would have two mainframe processors, each with a processing capacity of 5,900 KOPS [LIAS80, p. 104]. This configuration will allow for growth in processing capacity by upgrading the processors since even an average sized mainframe processor by current standards would have more processing capacity than the current 9020 systems. That is, a CCC has a total processor capacity of 790 KOPS and 1452 KOPS [LIAS80, p. 104] for 9020A and 9020D configurations, respectively.

Note that the usable capacity of a CCC is about 25% less than the total capacity due to memory contention and program element queueing. Mainframe processors with 7 to 9 times the capacity of a 9020D system [IBM80 and AMDA80] have been announced for delivery in 1982. By the end of the eighties, System/360 instruction-compatible mainframe processors with 25 to 50 times the capacity of a 9020D system [WISE80] should be available. By comparison, the estimates for the processing requirements for a fully automated ATC system are 10 to 15 times the processing capacity of a 9020D system [CLAP79].

The memory size of the current 9020 system is limited by signal propagation delay problems to 3 megabytes for a 9020A and 5 megabytes for a 9020D or 9020E. Current technology memory (solid state instead of core) is physcially much smaller in size than the 9020 memory size so that signal delays are no longer a problem. In addition, current technology memory is faster, cheaper and more reliable. Up to 32 megabytes of physical memory can be attached to candidate rehost computers. Even larger physical memories are possible since the near-term limit is based on effective usage rather than physical constraints.

The current size of the NAS monitor and application software is about 4.1 megabytes [FAAT81]; the size of the 9020E software is less than 1 megabyte. These programs in combination with a virtual machine monitor of 1.5 to 2 megabytes gives a minimum memory requirement for the rehost configuration of 6.6 to 7.1 megabytes. The expected memory requirements of near term ATC enhancements is less than 8 megabytes.

Mass storage devices (disk and tape) continue to improve with respect to capacity, response time and reliability. For example, disks with 200 to 500 megabyte capacities are currently in wide use; the capacity of a 2314 disk used with the 9020 system is 25 megabytes. Disks now on the market have about one half the access and latency delays of 2314 disks and about 3 times the data transfer rate of 2314 disks. At the present time, about 15 megabytes of information is stored on the 2314 disks in support of ATC operations [KAND77]. While the amount of disk-resident information is not expected to increase significantly, the capacity of a single replacement disk would allow at least 10 times more information to be disk-resident. Note that the large physical memory of the baseline rehost configuration would eliminate the need for disk buffering of program elements (PE's) which account for about 50% of disk activity [KAND77].

Tape drives with 8 times the capacity and 9 times the transfer rate of 2401 tapes are currently in use. Since data logging (SAR, REMON and DLOG) represent all of the tape usage, capacity or response time needs are not expected to change from the current situation. In addition, the tape drives would be assigned dedicated channels in the baseline rehost configuration so that channel utilization or contention issues in the current system [KAND77] would be minimized.

The above comments and the discussion in Sec. 3 indicate that the baseline rehost configuration will support growth in air traffic and ATC functions using existing hardware components that have significant performance characteristics (4 times the processing capacity of a 9020D system) but not maximal characteristics by current commercial standards. Since the NAS software can be rehosted without modifying the computer hardware, the processing capacity of the rehost computer can be upgraded. For example, the Amdahl 470/V7 can be field-upgraded to a model V8 and provide an increase in processor capacity that is estimated to be 7 percent by one source [LIAS80, p. 104] and 23 percent by another [HENK81, p. 14]. The survivability of 360 instruction-compatible processors is assured due to the very large investment in software for this type of processor. The impact of the transition to each upgrade option on the ATC operations would be minimal since the 360 instruction-compatible computers and peripherals have become de facto standards and market forces ensure that only fully compatible devices are offered for sale.

8.3 Software Growth Potential

While the growth potential of the hardware has been described in largely quantitative terms, the software growth potential is difficult to quantify and will be described in qualitative terms. Software evolution can proceed in two ways -- in a gradual, incremental extension of the rehosted NAS software or in a discontinuous replacement of the rehosted NAS software. In either case, the overall software organization for the baseline rehost configuration would be based on a virtual machine concept that allows many processes to proceed concurrently and as independently as necessary. The virtual macine concept would need an extension to allow efficient and responsive communications between cooperating subsystems operating in different virtual processes.

In the incremental extension case, the current NAS application software and an adapted NAS monitor would provide the kernel for developing new ATC processes. However, interfacing new processes with or revising and augmenting the existing NAS application software will continue to be a difficult task due to the strong data coupling between software modules and the highly optimized assembly language code used for some software modules. While rehosting the NAS application software will ensure ATC functional continuity, this rehosting will also preserve all the maintenance costs and problems of this software.

The second way that the software might change is through a rewrite and replacement of all (or at least a significant portion of) the NAS code. The idea is that after rehosting has been adopted and has taken care of the short-run capacity problems, longer run problems can possibly be dealt with by using modern software engineering techniques to develop new software that will eliminate the disadvantages of the current software and take advantage of the capabilities of the new hardware. This new software would not only provide more reliable and maintainable software support for the current ATC functions but would provide a more viable baseline for supporting the evolution of ATC functions. Note that cost for a software rewrite will be substantial and this cost has neither been estimated for the purposes of this report nor included in any cost calculations in Chapter 5.

REFERENCES

[AMDA80] Amdahl Product Announcement for Amdahl 580 Series, 1980.

[CLAP79] David Clapp, Joseph Hagopian and Ronald Rutledge, Analysis of Expandability and Modifiability of Computer Configuration Concepts for ATC, Volume I, FAA-EM-79-22, November 1979.

[FAAT81] FAA Technical Center, informal discussions with ACT-700 staff, Atlantic City, May 1981.

[HENK81] Tom Henkel, "IBM and the Plug-Compatibles," Computerworld, July 13, 1981, pp. 12-16.

[IBM80] IBM Product Announcement for IBM 3081 Processors, 1980.

[KAND77] William Kandler, Gary Nielsen and John Squires, Response Time Analysis Study Fort Worth ARTCC Measurement, Interim Report, Logicon R4940-107, August 1977.

[LIAS80] Edward J. Lias, "Tracking the Elusive KOPS," Datamation, November 1980, pp. 99-105.

[WISE80] Kensall Wise, Kan Chen and Ronald Yokely, Microcomputers: A Technology Forecast and Assessment to the Year 2000, John Wiley and Sons, 1980.

A. SYSTEM AVAILABILITY AND SYSTEM MTBF: DETAILED ANALYSIS

A.1 Purpose and Organization of this Appendix

The purpose of this appendix is to show how the estimates of system availability and system mean time between failure (MTBF) given in Sec.'s 2.2 and 2.5 are obtained. First, considering only hardware failures, Sec. A.2 and A.3 explain the principles used to estimate system availability and MTBF, respectively. Second, Sec. A.4 extends the analysis to cover not only hardware failures but also software failures. The principles in each section are illustrated by calculating the availabilities and MTBF's for the 9020D/9020E system and for the rehost system under the baseline assumptions.

A.2 System Availability: Hardware

A.2.1 Principles

This subsection shows how system availability can be estimated from information about failure rates, repair rates, and configurations.

Consider a single unit with an MTBF of $1/\lambda$. Let the mean time to repair be $1/\mu$. Define a "cycle" to start at the moment a unit is placed in service following a repair and to last until the next moment when, having failed and been repaired, the unit is again placed in service. A cycle, therefore, includes the time spent operating and the time spent being repaired. The expected time spent operating in a cycle is $1/\lambda$, and the expected time spent being repaired is $1/\mu$, so the expected length of the cycle is

$$\frac{1}{\lambda} + \frac{1}{\mu} = \frac{\mu + \lambda}{\lambda\mu} .$$

The fraction of the time spent operating is

$$\frac{1/\lambda}{(\mu + \lambda)/\lambda\mu} = \frac{\mu}{\mu + \lambda} .$$

Therefore, $\mu/(\mu + \lambda)$ is this single unit's availability, i.e., the probability that it is operating at any randomly chosen point in time.

Consider a subsystem that contains n identical, independent units; suppose that at least m of the units must be working to prevent a failure of the subsystem. Letting A(n,m) stand for the availability of this subsystem of n units, m of which must operate to prevent a failure, we have

$$A(n,m) = \sum_{i=m}^{n} \left[C(n,i) \left(\frac{\mu}{\mu + \lambda}\right)^{i} \left(\frac{\lambda}{\mu + \lambda}\right)^{n-i}\right], \tag{1}$$

where

$$C(n,i) = \frac{n!}{i!(n-i)!},$$

which is the number of different combinations of i objects that can be chosen from a set of n objects. The values of the availability function that are used below are

$$\begin{aligned} A(2,1) &= 2 \left(\frac{\mu}{\mu + \lambda}\right) \left(\frac{\lambda}{\mu + \lambda}\right) + \left(\frac{\mu}{\mu + \lambda}\right)^2 \\ &= \frac{2\mu\lambda + \mu^2}{(\mu + \lambda)^2} \\ &= 1 - \left(\frac{\lambda}{\mu + \lambda}\right)^2 . \\ A(3,2) &= 3 \left(\frac{\mu}{\mu + \lambda}\right)^2 \left(\frac{\lambda}{\mu + \lambda}\right) + \left(\frac{\mu}{\mu + \lambda}\right)^3 \\ &= \frac{3\mu^2\lambda + \mu^3}{(\mu + \lambda)^3} . \\ A(6,5) &= 6 \left(\frac{\mu}{\mu + \lambda}\right)^5 \left(\frac{\lambda}{\mu + \lambda}\right) + \left(\frac{\mu}{\mu + \lambda}\right)^6 \\ &= \frac{6\mu^5\lambda + \mu^6}{(\mu + \lambda)^6} . \end{aligned}$$

These formulas allow the availability of any subsystem to be estimated once λ and μ are known.

It will now be explained how the availability of a system is built up from the availability of its subsystems. If a system is composed of a number of independent subsystems such that a failure of any subsystem causes a system failure, then the system availability A_s is the product of the subsystem availabilities. For example, if there are two independent subsystems with availabilities A_1 and A_2, then

$$A_s = A_1A_2. \tag{2}$$

If a system is composed of two independent subsystems and if a system failure occurs only if both subsystems fail, then the system availability is the probability that at least one of the subsystems works, i.e.,

$$\begin{aligned} A_s &= 1 - (1-A_1)(1-A_2) \\ &= A_1 + A_2 - A_1A_2. \end{aligned} \tag{3}$$

Equations (2) and (3) are only valid if no two subsystems are both failed at the same time. While this might happen, in the problem being considered it has such a small probability (on the order of about 10^{-12}) that it can be ignored without damaging the results.

A.2.2 Results

<u>9020D/9020E system.</u> The availability of a system with a 9020D in the CCC and a 9020E in the display channel will now be estimated using the baseline component MTBF's from Table 2-1 and a mean time to repair (MTTR) of one hour.

Table A-1 shows the details of the calculation of the availability of a 9020D system. It shows for each subsystem the number of units n it has, the number m that must be working to prevent a system failure, the failure rate (from Table 2-1), and the relevant availability formula from Sec. A.2, where it is assumed that the mean time to repair $1/\mu$ is 1 hour. The last column in the table shows the subsystem availability, which is obtained by substituting the λ given in the table into the formula. Since the failure of any subsystem causes a system failure of the 9020D, the availability of the 9020D is 0.99999151, which is the product of the subsystem availabilities.

TABLE A-1: AVAILABILITY OF A 9020D SYSTEM

Component	n	m	λ	Formula	Availability
CE	3	2	$\frac{1}{2301}$	$\frac{3\lambda + 1}{(1 + \lambda)^3}$	0.99999943
SE	6	5	$\frac{1}{3137}$	$\frac{6\lambda + 1}{(1 + \lambda)^6}$	0.99999851
IOCE	3	2	$\frac{1}{3161}$	$\frac{3\lambda + 1}{(1 + \lambda)^3}$	0.99999970
TCU	3	2	$\frac{1}{8482}$	$\frac{3\lambda + 1}{(1 + \lambda)^3}$	0.99999996
SCU	3	2	$\frac{1}{700}$	$\frac{3\lambda + 1}{(1 + \lambda)^3}$	0.99999390
9020D System					0.99999151

Since the 9020D/9020E system fails if the 9020D or the 9020E fails, the availability of the 9020D/9020E system is the availability of the 9020D times that of the 9020E. In this report it is assumed that the 9020D and 9020E are equivalent from the point of view of reliability. Therefore, the availability of a system with a 9020D in the CCC and a 9020E in the display channel is 0.99998301, which is the square of 0.99999151.

Rehost system. The availability of the rehost system will now be estimated. Define a mainframe to be a CPU, memory, and six pairs of channels. Since the two mainframes run in parallel, it is first necessary to find the availability of a single mainframe.

Table A-2 shows the details of the calculations of the availability of a single mainframe. This table shows that for a single mainframe to be available, the single CPU must be working, the single memory must be working, and at least one channel in each of the six pairs must be working. This table also shows the failure rates (from Table 2-1) and the formulas used to calculate availability (from Sec. A.2). The last column shows the availability of each subsystem, and by multiplying these three availabilities together the availability of a single mainframe is found to be 0.99751951.

Since a system failure occurs only if both mainframes fail, we are interested in the probability that at least one mainframe is working, which from Eq. (3) is

$$
\begin{aligned}
1 - (1-A)(1-A) &= 2A - A^2 \\
&= 2(0.99751943) - (0.99751943)^2 \\
&= 0.99999385.
\end{aligned}
$$

Table A-3 shows the details of the calculations of the availability of the rehost system. At least 1 of the 2 mainframes, TCU's, and SCU's must be working to prevent a system failure. The failure rates and availability formulas are given for the TCU's and SCU's (but not for the mainframe since the special calculation above replaces the formula). The last column shows the availability of each subsystem, and the availability of the rehost system is 0.99999283, the product of these availabilities.

TABLE A-2: AVAILABILITY OF A SINGLE MAINFRAME

Component	n	m	λ	Formula	Availability
CPU	1	1	$\frac{1}{1356}$	$\frac{1}{1+\lambda}$	0.99926308
Memory	1	1	$\frac{4}{2293}$	$\frac{1}{1+\lambda}$	0.99825860
Channels	2	1	$\frac{1}{1316}$	$[1-(\frac{\lambda}{1+\lambda})^2]^6$	0.99999654
Mainframe					0.99751951

TABLE A-3: AVAILABILITY OF THE REHOST SYSTEM

Component	n	m	λ	Formula	Availability
Mainframe	2	1			0.99999385
TCU	2	1	$\frac{1}{1000}$	$1-(\frac{\lambda}{1+\lambda})^2$	0.99999900
SCU	2	1	$\frac{1}{7163}$	$1-(\frac{\lambda}{1+\lambda})^2$	0.99999998
System					0.99999283

A.3 System MTBF: Hardware

A.3.1 Principles

This subsection shows how system MTBF can be estimated from information that is known.

The definition of availability is

$$A_s = \frac{MTBF_s}{MTBF_s + MTTR_s} ,$$

where the subscript s is used to show we are talking about the system. This can be rearranged to

$$MTBF_s = \frac{A_s}{1-A_s} MTTR_s . \tag{4}$$

Eq. (4) expresses the system MTBF as a function of system availability and the system MTTR. The values for system availability are known; they are derived in Sec. A.2. Therefore, once the system MTTR is known, system MTBF can be estimated. Note that system MTTR is the same thing as the expected duration of a system failure. The method used to estimate the system MTTR will now be explained.

One might think that because the MTTR for each unit is 1 hour, the MTTR for the system would also be 1 hour, but this is incorrect. To see this, suppose that one unit fails; this does not cause a system failure because of redundancy. Then suppose that a second unit fails; this does cause a system failure. Since the repair times are distributed exponentially, and hence memoryless, we can let t=0 be the time of the second failure. If the system failure were caused by one unit failing and if the failure ended when that unit was repaired, then the system MTTR would be 1 hour. However, since 2 units have failed and since the system failure ends when either unit is repaired, the expected duration of the MTTR is 1/2 hour. That is, if t_1 is the time that the first unit is repaired, and t_2 is the time that the second is repaired, then the duration of the failure is $\min(t_1,t_2)$, and the expectation of $\min(t_1,t_2)$ is 1/2. This is an implication of the following theorem.

<u>Theorem</u>: If $t_1,\ldots, t_n$ are independent, exponential random variables with means $1/\mu_1,\ldots,1/\mu_n$, respectively, then the mean of min $(t_1,\ldots,t_n)$ is $1/(\mu_1+\ldots+\mu_n)$.

<u>Proof</u>:

Since the n random variables $t_1, \ldots , t_n$ are independent, for any t

$$\Pr[\min(t_1, \ldots , t_n)>t] = \Pr[t_1>t] \times \ldots \times \Pr[t_n>t]$$

$$= e^{-\mu_1 t} \times \ldots \times e^{-\mu_2 t}$$

$$= e^{-(\mu_1+\ldots+\mu_n)t}.$$

The last equality implies that the distribution function of $[\min(t_1,\ldots,t_n)<t]$ is $1-e^{-(\mu_1+\ldots+\mu_n)t}$, which is the distribution function of an exponential distribution with mean $1/(\mu_1+\ldots+\mu_n)$. This completes the proof.

In the baseline case we assume that the unit MTTR is 1 hour, i.e., $\mu =$ 1. This theorem then implies that:

- the system MTTR is 1/2 hour if the system failure is caused by the failure of two units (e.g., by 2 9020D CE's or by two rehost TCU's);

- the system MTTR is 1/3 hour if the system failure is caused by the failure of three units (e.g., by a pair of channels failing on one rehost mainframe and the CPU failing in the other rehost mainframe);

- the system MTTR is 1/4 hour if the system failure is caused by the failure of four units (e.g., by a pair of channels failing on each rehost mainframe).

The MTTR for each system is calculated in the next subsection.

A.3.2 Results

90200/9020E system. The MTBF of a system with a 9020D in the CCC and a 9020E in the display channel will now be estimated. Subsec. A.2.2 estimates the availability of the 9020D/9020E system to be 0.99998301. Since this system fails if and only if two like units fail, the theorem of Subsec. A.3.1 implies that the system MTTR is 1/2. Eq. (4) can now be used and the system MTBF is

$$\begin{aligned} MTBF_s &= \frac{A_s}{1-A_s} MTTR_s \\ &= \frac{0.99998301}{1-0.99998301} \times \frac{1}{2} \\ &= 29{,}429 \text{ hours} \\ &= 1226 \text{ days.} \end{aligned}$$

Therefore, the estimate of the MTBF of the system with a 9020D in the CCC and a 9020E in the display channel is 1226 days. Each system outage is expected to last a half hour.

Rehost system. The MTBF of the rehost system will now be estimated. The availability of the rehost system is estimated in Subsec. A.2.2 to be 0.99999283. It is claimed that the system MTTR does not differ significantly from 1/2; this claim is substantiated below. The system MTBF can now be estimated using eq. (4).

$$\begin{aligned} MTBF_s &= \frac{A_s}{1-A_s} MTTR_s \\ &= \frac{0.99999283}{1-0.99999283} \times \frac{1}{2} \end{aligned}$$

$= 69,735$ hours

$= 2905$ days.

Therefore, the MTBF of the rehost system is estimated to be 2905 days. Each system outage is expected to last a half hour.

All that remains is to substantiate the claim that the MTTR for the rehost system does not differ significantly from 1/2. The reason why the rehost system MTTR is not exactly 1/2 is because, unlike the 9020D/E system, not all system failures are caused by two units failing. There are three cases. First, a system failure might result from the failure of two units, i.e., both SCU's fail, both TCU's fail, both CPU's fail, both memories fail, or a CPU fails on one mainframe and a memory fails on the other. In this case the system MTTR is 1/2 hour. Second, a system failure might result from the failure of three units, i.e., a pair of channels fails on one mainframe and the CPU or memory fails on the other. In this case the system MTTR is 1/3 hour. Third, a system failure might result from the failure of four units, i.e., a pair of channels fails on each mainframe. In this case the system MTTR is 1/4 hour.

The system MTTR, therefore, is

$$MTTR_S = (1/2)p_2 + (1/3)p_3 + (1/4)p_4, \qquad (5)$$

where p_i is the probability that the system failure is caused by the failure of i units, given that a system failure occurs. To show that the rehost system MTTR does not differ significantly from 1/2, the weights p_2, p_3, and p_4 will now be found. Table A-4 shows the relevant information. The first column shows the ways that a system failure can occur. For example, "CPU x CPU" means that the CPU's fail in both mainframes. Since this can only happen in one way, a 1 is written in the second column. "CPU x Mem." means that the CPU fails in one mainframe and the memory in the other; this can happen in two ways since the CPU can fail in either mainframe. "CPU x Ch." means that the CPU fails in one main frame and a pair of channels fails in the other. Since each main frame has six pairs of channels, there are 12 different ways that "CPU x Ch." can occur.

TABLE A-4: PROBABILITIES OF DIFFERENT TYPES OF SYSTEM FAILURES IN THE REHOST SYSTEM

Two unit failures	Number of Combinations	Unconditional Probability	Conditional Probability
CPU x CPU	1	5.431×10^{-7}	0.0757
CPU x Mem.	2	2.567×10^{-6}	0.3576
Mem. x Mem.	1	3.032×10^{-6}	0.4223
TCU	1	1.000×10^{-6}	0.1393
SCU	1	2.000×10^{-8}	0.0028
Subtotal			0.9977
Three unit failures			
CPU x Ch.	12	5.129×10^{-9}	0.0007
Mem. x Ch.	12	1.212×10^{-8}	0.0017
Subtotal			0.0024
Four unit failures			
Ch. x Ch.	36	1.211×10^{-11}	0.0000
Total		7.179×10^{-6}	1.0001

The rest of the failures shown in Table E-4 have similar explanations except for the TCU and SCU; TCU, for example, means that both tape units fail.

The availability for a single CPU is 0.99926308, for a single memory is 0.9982586, and for a single pair of channels is 0.99999942; these figures are from Table A-2 (where the channel availability is obtained by taking the sixth root of the number shown). To illustrate the method of calculating the unconditional probability of any particular type of mainframe failure, consider the CPU x Mem. failure. Since the event of a CPU failing is independent of the event of a memory failing, and since there are two different ways a CPU x Mem. failure can occur, the probability of a CPU x Mem. failure is

$$(1 - 0.99926308)(1 - 0.99825860)2 = 2.567 \times 10^{-6}.$$

The interpretation of this probability is that if a point in time is chosen at random, then the probability of a CPU x Mem. failure obtaining at that point is 2.567×10^{-6}. The other unconditional probabilities are similarly calculated and entered in Table A-4. The exception is that the TCU and SCU unconditional probabilities are merely one minus the availabilities in Table A-3. The sum of the probabilities is 7.179×10^{-6}, which is the probability of the system being down. (The availability of the two mainframes is then

$$1-7.179 \times 10^{-6} = 0.99999282,$$

which checks with the number in Table A-3 derived by a different method.)

Dividing each unconditional probability in Table E-4 by 7.179×10^{-6} gives the conditional probabilities shown in the last column. For example, given that there is a system failure, the probability that this is a CPU x CPU failure is 0.0757.

The subtotals in Table A-4 give the values for the weights p_2, p_3, and p_4. Eq. (5) now becomes

$$MTTR_s = (1/2)p_2 + (1/3)p_3 + (1/4)p_4$$

$$= (1/2)(0.9977) + (1/3)(0.0024) + (1/4)0$$
$$= 0.4998.$$

It is seen that the system MTTR, rounded to three significant figures, equals 1/2. This completes the argument that the rehost system MTTR does not differ significantly from 1/2.

A.4 System Availability and System MTBF: Hardware and Software

Sec. 2.5 contains estimates of the system availability and MTBF that take into account not only hardware failures but also software failures. This section shows how these estimates are obtained. That is, this section shows how the methods of Sec.'s A.2 and A.3, which only dealt with hardware, can be extended to cover the case of both hardware and software. Again the calculations will be illustrated using the baseline assumptions. The 9020D/9020E system and the rehost system are considered separately.

9020D/9020E system. The analysis proceeds in five steps. First, estimate the mean duration of a software outage. It is assumed in Sec. 2.5 that, given that there is a software failure, the system outage is 0.5 minute with probability 0.9 and is 15 minutes with probability 0.1. The mean outage due to a software failure is then

$$(0.5 \times 0.9) + (15 \times 0.1) = 1.95 \text{ min.}$$

Second, estimate the software availability. Sec. 2.5 assumes that the NAS software has the same MTBF as the 9020D/9020E hardware; this MTBF is 1226 days in the baseline case. Since 1226 days contains 1,765,440 minutes, the software availability is

$$\frac{1,765,440}{1,765,440 + 1.95} = 0.99999890.$$

Third, estimate system availability. Since the system works only if both the hardware and the software work and since the hardware availability is 0.99998301, eq. (2) implies that the system availability considering both hardware and software is

$$0.99998301 \times 0.99999890 = 0.99998191.$$

Fourth, estimate the system MTTR, i.e., the mean duration of a system outage. Since it is assumed that the number of system failures caused by hardware equals that caused by software, since the mean duration of a hardware outage is a half hour, and since the mean duration of a software outage is 1.95 minutes, the mean duration of a system outage is

$$\begin{aligned} MTTR_s &= (1/2 \times 1/2 \text{ hr}) + (1/2 \times 1.95 \text{ min} \times 1 \text{ hr}/60 \text{ min}) \\ &= 0.26625 \text{ hr}. \end{aligned}$$

Fifth, estimate system MTBF using eq. (4).

$$\begin{aligned} MTBF_s &= \frac{A_s}{1-A_s} MTTR_s \\ &= \frac{0.99998191}{1-0.99998191} \times 0.26625 \\ &= 14718 \text{ hr} \\ &= 613 \text{ days}. \end{aligned}$$

Rehost system. The analysis of Sec.'s A.2 and A.3 must be changed in two ways -- to include failures in the virtual machine monitor (VMM) and in the NAS software. VMM will be treated as another component of a mainframe just like a CPU or a memory; this is because if the VMM in one mainframe fails, then processing is switched to the other mainframe, and the first mainframe is out while being restarted. The analysis proceeds in six steps.

First, estimate the availability of the VMM. With a failure rate of twice per month, the VMM runs on average for 15 days (360 hours) and is then down for 1/6 hour. The VMM availability then is

$$\frac{360}{360 + 1/6} = 0.99953725.$$

Second, estimate the availability of a single mainframe (including the VMM). Since the CPU, memory, channels, and VMM must all be working if the mainframe is to work, the mainframe availability is the availability of the first three (which from Table A-2 is 0.99751951) times the availability of the VMM, or

$$0.99751951 \times 0.99953725 = 0.99705791.$$

Third, estimate the availability of the two mainframes, i.e., the probability that at least one mainframe is working. Eq. (3) implies

$$2(0.99705791) - (0.9970591)^2 = 0.99999134.$$

Fourth, estimate the availability of the rehost system. It is the product of the availabilities of the mainframes, the TCU, the SCU's, and the NAS software. We have

$$0.99999134 \times 0.99999900 \times 0.99999998 \times 0.99999890 = 0.99998922.$$

Fifth, estimate the mean duration of a system outage. Table A-5, which is analogous to Table A-4, shows the different types of system failure, each one's mean time to repair (which assumes that the repair times are exponentially distributed and uses the theorem in A.3.1), and the probability of each type of failure given that there is a system failure. These conditional probabilities are used as weights in taking a weighted average of the mean repair times, and the resulting system MTTR is

$$(1/2 \times 0.6628) + (1/3 \times 0.0016) + (1/4 \times 0) + (1/7 \times 0.2123) + (1/8 \times 0.0018) + (1/12 \times 0.0198) + (1.95/60 \times 0.1018) = 0.3674 \text{ hr.}$$

Sixth, estimate the system MTBF using eq. (4).

$$\begin{aligned} MTBF_s &= \frac{A_s}{1-A_s} MTTR_s \\ &= \frac{0.99998922}{1-0.99998922} \times 0.3674 \\ &= 34{,}081 \text{ hours} \\ &= 1420 \text{ days.} \end{aligned}$$

TABLE A-5: PROBABILITIES OF DIFFERENT TYPES OF SYSTEM FAILURES IN THE REHOST SYSTEM

MTTR	Type of Failure	Number of Combinations	Unconditional Probability	Conditional Probability
1/2	CPU x CPU	1	5.431×10^{-7}	0.0503
1/2	CPU x Mem.	2	2.567×10^{-6}	0.2375
1/2	Mem x Mem.	1	3.032×10^{-6}	0.2806
1/2	TCU	1	1.000×10^{-6}	0.0925
1/2	SCU	1	2.000×10^{-8}	0.0019
	Subtotal			0.6628
1/3	CPU x Ch.	12	5.129×10^{-9}	0.0005
1/3	Mem. x Ch.	12	1.212×10^{-8}	0.0011
	Subtotal			0.0016
1/4	Ch. x Ch.	36	1.211×10^{-11}	0.0000
1/7	CPU x VM	2	6.820×10^{-7}	0.0631
1/7	Mem. x VM	2	1.612×10^{-6}	0.1492
	Subtotal			0.2123
1/8	VM x Ch.	12	1.932×10^{-8}	0.0018
1/12	VM x VM	1	2.141×10^{-7}	0.0198
1.95/60	NAS Software	1	1.100×10^{-6}	0.1018
	Total		1.081×10^{-5}	1.0001

B. SYSTEM MTBF WITHOUT REPAIRS: DETAILED ANALYSIS

B.1 Purpose and Organizaton of this Appendix

The purpose of this appendix is to show how the estimates in Sec. 2.3 of system MTBF are obtained under the assumption that no repairs are made. Sec. B.2 derives the equations for the system with a 9020D in the CCC and a 9020E in the display channel. Sec. B.3 then derives the equations for the rehost system. Sec. B.4 describes the approximation used to obtain the system MTBF. This appendix only considers hardware failures.

Throughout this appendix reliability is defined to mean the probability that there has not been a failure after a stated period of operation.

B.2 Reliability of the 9020D or 9020E Configuration

The 9020D system is working if all of the following five conditions hold:

- at least 2 of the 3 CE's are working;
- at least 5 of the 6 SE's are working;
- at least 2 of the 3 IOCE's are working;
- at least 2 of the 3 TCU's are working;
- at least 2 of the 3 SCU's are working.

It is assumed that the 9020E has the same reliability as the 9020D since these two systems have the same configuration (except that in the 9020E some of the SE's are replaced by display elements).

The function $r(t)$ is used to denote reliability for a single component, e.g., individual CE's, SE's, IOCE's, TCU's, and SCU's. The function $R(t)$ is used to denote reliability for a subsystem or system. Subscripts are used to distinquish different reliability functions. For example, $r_{CE}(t)$ is the probability that a single, specified CE has not failed after t hours of operation; $R_{CE}(t)$ is the probability that at most 1 of the 3 CE's have failed.

Exponential failure rates are assumed for each component. Therefore, for any component

$$r(t) = e^{-\lambda t} \tag{1}$$

where λ is the failure rate.

For the subsystems containing three components with one redundant (i.e., the CE, IOCE, TCU, or SCU), that subsystem will function if all three components function or if any two components function. Mathematically,

$$R_{CE} = r^3 + 3r^2(1-r), \tag{2a}$$

$$R_{IOCE} = r^3 + 3r^2(1-r), \tag{2b}$$

$$R_{TCU} = r^3 + 3r^2(1-r), \text{ and} \tag{2c}$$

$$R_{SCU} = r^3 + 3r^2(1-r), \tag{2d}$$

For clarity, the subscripts have been dropped on the right hand side and the t's have been dropped on both sides. Eq. (2a) can be explained in the following way. R_{CE} is the probability that not more than one CE will fail. This probability is calculated by substituting into (2a) the value of r obtained by substituting into eq. (1) the relevant failure rate, i.e., λ_{CE}.

For the storage elements where we have a total of six elements, the appropriate expression is

$$R_{SE} = r^6 + 6r^5(1-r). \tag{3}$$

The overall system reliability for the 90200 system is found by combining equations (2) and (3) to obtain

$$\begin{aligned} R_{90200} = {} & [r_{CE}^3 + 3r_{CE}^2(1-r_{CE})]\,[r_{SE}^6 + 6r_{SE}^5(1-r_{SE})]\,[r_{TCU}^3 + 3r_{TCU}^2(1-r_{TCU})] \\ & [r_{SCU}^3 + 3r_{SCU}^2(1-r_{SCU})]\,[r_{IOCE}^3 + 3r_{IOCE}^2\,(1-r_{IOCE})]. \end{aligned} \tag{4}$$

That is, the system reliability is the product of the probabilities that no more than one component failure occurs in any of the subsystems.

The reliability of a system with a 9020D in the CCC and a 9020D in the display channel is the product of the reliability of the 9020D and the reliability of the 9020E, i.e.,

$$R_{9020D/9020E}(t) = [R_{9020D}(t)]^2. \tag{5}$$

B.3 Reliability of the Rehost Configuration

The rehost system contains:

- two mainframes, where a mainframe is defined to include a CPU, an SE, and twelve channels;
- two TCU's;
- two SCU's.

The rehost system is operating successfully if all three of the following conditions are satisfied:

- at least one mainframe is operating;
- at least one TCU is operating; and
- at least one SCU is operating.

Since a mainframe operates only if its CPU, its SE, and at least one channel in each of the six pairs operates, the reliability of a single mainframe is given by

$$R_M = R_{CPU} R_{SE} R_{CH}. \tag{6}$$

Since there is only one CPU and one SE in the mainframe, these components have simple exponential failure rates. The reliability for the six pairs of channels is

$$R_{CH} = [r^2 + 2r(1-r)]^6. \tag{7}$$

Therefore, eq. (6) now implies that the reliability for a single mainframe is

$$R_M = r_{CPU} \, r_{SE} \, [r_{CH}^2 + 2r_{CH} \, (1 - r_{CH})]^6. \tag{8}$$

For the TCU's and SCU's, one of two must function, so the reliabilities of these subsystems are

$$R_{TCU} = r_{TCU}^2 + 2r_{TCU}(1-r_{TCU}), \text{ and} \tag{9}$$

$$R_{SCU} = r_{SCU}^2 + 2r_{SCU}(1-r_{SCU}). \tag{10}$$

The reliability of the complete rehost system is the

$$R_S = [R_M^2 + 2R_M(1-R_M)] \, R_{TCU} \, R_{SCU} \tag{11}$$

where the complete expression is obtained by substituting from (8), (9), and (10) into (11).

B.4 Approximating System MTBF

The reliability as a function of time is given for the 9020D/9020E system by eq. (5) and for the rehost system by eq. (11). The method used to extract from these functions the system MTBF's shown in Tables 2-5 and 2-6 is as follows. For a truly exponential reliability function, the MTBF is the time at which reliability is equal to 0.37 (i.e., to 1/e). In the present case, even though each unit has an exponential reliability function, because of redundancy the system reliability function is not exponential. Nevertheless, the time at which reliability equals 0.37 is used to approximate the system MTBF. Therefore, the times shown in Tables 2-5 and 2-6 are only approximate MTBF's; strictly speaking, these are the times that elapse between when the system begins running and when the reliability drops to 0.37.

C. PARTIAL REPLACEMENT

The body of this report has assumed that if rehosting is adopted, then the 9020's will be replaced at all twenty ARTCC's. However, since hardware is the major cost (see Sec. 5.7), it has been suggested that the 9020's only be replaced at those centers that face capacity problems; proponents of this idea claim that this partial replacement would take care of the capacity problems while minimizing the cost. The purpose of this appendix is to point out the advantages and disadvantages of partial replacement.

There are three main disadvantages to partial replacement. First, with two entirely different systems in the field, support would be greatly complicated since training at the Academy, inventory management at the Depot, and support at the Technical Center would need to be carried out for the different systems. Second, since both the level of air traffic and the lifetime of the rehost system are hard to predict, there is some doubt as to exactly which centers will face capacity problems and which will, therefore, require a new system. Third, if the view is taken that the rehost system will evolve into a full replacement system (see Ch. 8), then money will probably not be saved by partial replacement. The expenditure on new hardware could not be eliminated; it could only be delayed.

The main advantage of partial replacement is the cost saving. An additional advantage is that the number of transitions could be held down. Some of the relevant costs will be estimated, but it should be stressed that some costs cannot be quantified, e.g., costs due to the inconvenience or confusion of having multiple systems. Therefore, the discussion below should be thought of as a treatment of some of the costs and not as a complete treatment. Six areas in which the cost of partial replacement differs from that of total replacement will now be discussed. For concreteness, assume that Amdahl 470/V7's are the mainframes that are procured. Consider the cases where replacement is at 5 and 10 ARTCC's.

First, since new systems need not be purchased for the centers which retain the 9020's, there is a saving in the hardware acquisition cost. Since this cost is $5.4 million per center (see Table 5-9), the total saving compared to total replacement would be $81.0 million if there is replacement at 5 centers and $54 million if there is replacement at 10.

Second, an initial inventory of new spare parts need not be laid in for the centers which retain the 9020's. Since the stock of spares is estimated to be $1.3 million per center (see Subsec. 5.4.3), the saving relative to total replacement is $19.5 million or $13.0 million if replacement were at 5 or 10 centers, respectively.

Third, no transition cost would be incurred at centers at which there is no replacement. The per center saving is $1 million on remodeling cost, $200,000 in extra personnel cost, and $370,000 for training cost (see Sec. 5.5). The saving in transition cost is then $1.6 million per center, for a total of $24.0 million or $16 million if replacement is at 5 or 10 centers, respectively.

Fourth, since the procurement would not take as long if there were replacement at fewer centers, there will be a saving in the program management and support cost. The procurement would be shortened by 18 or 12 months if there were replacement at 5 or 10 centers, respectively. Since the program management is expected to cost $6.3 million per year during deployment (see Table 5-19), this means the cost saving would be $9.5 million or $6.3 million, respectively, for the two cases.

Fifth, if there is no replacement at a center, then it does not reap the annual saving in maintenance cost provided by the more reliable new system. After the three-year shakedown period, the saving in maintenance personnel cost is $412,292 per center (see Table 5-10). Therefore, the annual cost penalty of not replacing is $6.2 million or $4.1 million, depending on whether replacement is at 5 or 10 centers. With full replacement, the annual parts saving is $1,013,000 million; $0.8 or $0.5 million of this would not be saved if replacement were only at 5 or 10 centers. Thus, the total annual cost penalty paid is $7.0 million if there is replacement at 5 centers or $4.6 million if there is replacement at 10.

Sixth, extra personnel would be required at the Technical Center since two systems would need support. Table C-1 shows the number of additional people that, it is estimated, would be required at the Technical Center if there were partial replacement. The total annual cost of these extra personnel is rounded to $3 million. (This table assumes that the average grade is GS-13, step 4, in AF; it is halfway between GS-13, step 4, and GS 14, step 4, in AT and R&D; and it is GS-11 Step 4, in ACT-700. The salaries for GS-11, 13, and 14, Step 4, are currently $24,736, $35,252, and $41,657 respectively. To these salaries 10 percent has been added to cover benefits and 5 percent to cover overtime.)

These figures are summarized in Table C-2. Compared to full replacement, the saving in initial cost is $134.0 million if there is replacement at 5 centers and $89.3 million if there is replacement at 10. In other words, the initial cost is $107.0 million, $151.7 million, or $241.0 million depending on whether replacement is at 5, 10, or 20 centers. Compared to full replacement, the long-term increase in the annual cost is $10.1 million if there is replacement at 5 centers and $7.7 million if there is replacement at 10. In other words, with full replacement the saving in annual cost is $9.3 million; with replacement at 10 centers the saving is $1.6 million; with replacement at 5 centers the annual cost rises by $0.8 million.

TABLE C-1: INCREASED ANNUAL PERSONNEL COST AT THE TECHNICAL CENTER IF THERE IS PARTIAL REPLACEMENT

Organization	Number of New Personnel	Average Cost	Total Cost
AT	30	$44,223	$1,326,690
AF	20	40,540	810,800
R & D	10	44,223	442,230
ACT-700	20	28,446	568,920
Total			$3,148,640

TABLE C-2: AMOUNT SAVED IF THERE IS PARTIAL REPLACEMENT (millions)

	Number of Centers at which there is Replacement	
	5	10
One-Time Cost Saving		
Hardware acquisition	$ 81.0	$ 54.0
Initial parts inventory	19.5	13.0
Transition cost	24.0	16.0
Program management	9.5	6.3
Total	$134.0	$ 89.3
Annual Cost		
Maintenance Cost	($ 7.0)	($ 4.6)
Extra personnel at the FAATC	(3.1)	(3.1)
Total	($10.1)	($ 7.7)

N.B. A figure in parentheses denotes a cost increase rather than a cost saving.

D. WHY SPECIAL HARDWARE IS NEEDED IN THE REHOST SYSTEM

The rehost system as described in Sec. 1.3 contains two pieces of special hardware--the radar input line multiplexor (RIN/LM) and the display buffer. This appendix states why this special hardware is needed and what its function is.

Any rehost computer system must be capable of interfacing to the radar circuits and the display generators. The radar circuits provide raw data on the location and identity of controlled aircraft. The display generators maintain the geo-situation plot for the controller suites. In the current 9020 system, both of these interfaces are supported with special purpose extensions of the basic system. That is, the radar input processing program provides device support for the radar circuits and runs continually in an IOCE after it is dispatched during startup/startover. All other devices are supported in a traditional interrupt-driven manner. The display generators are supported with special access to display buffers in the display channel memory to ensure adequate response for the display refresh process. Each of these interfaces represents special problems for a rehost system.

RIN/LM. If the radar circuits were directly connected to a rehost computer system, then the mainframe would be required to provide support equivalent to the current RIN support. The alternatives for this radar input support are either to allow a channel program to run continually (equivalent to the current approach) or to use the traditional interrupt driven support. Either of these alternatives would severely compromise the performance of the mainframe due to the frequency and response requirements for radar data processing.

An alternative for supporting radar data input is to provide a pre-processor for each radar circuit that would perform the RIN function and present valid and reformatted radar data to the mainframe in blocked records so that one mainframe I/O operation would access several radar data values. This approach would ensure adequate capacity to process the raw radar data and avoid potential performance problems with the mainframe.

(It should be mentioned that Amdahl has studied the radar input problem and has tentatively decided that the RIN program can run in the mainframe without causing a performance problem. If this analysis proves correct, then the RIN/LM can be omitted from the rehost system.)

Display buffer. The display generators require periodic access to display buffers which define the geo-situation plot so that the plan view displays in the controller suites can be dynamically updated. In the current 9020 system, these display buffers are provided as part of the memory in the display channels which allows them to be updated by the Central Computer Complex (CCC) and to be accessed by the display generators. In the baseline rehost configuration, one large memory would serve both the CCC functions as well as the display functions. If the display buffers were to be resident in the mainframe memory, the frequency and response requirements for the display generator accesses to these buffers would significantly degrade the overall memory performance for the remainder of the system. An alternative approach for resolving this potential mainframe problem is to provide a capability between the mainframe and the display generators. That is, the display buffers would be maintained in special purpose memory units that could be loaded and updated by the mainframe and accessed, as necessary, by the display generators. Again, this approach would ensure adequate capacity to service the display generators and avoid potential mainframe performance problems.